REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
07-12-2011	Final Report	28-Jul-2010 - 27-Apr-2011

4. TITLE AND SUBTITLE

Summary Report for the Army Research Organization (ARO) Workshop on Social Trust Computing

5a. CONTRACT NUMBER

W911NF-10-1-0350

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

611102

6. AUTHORS

Jennifer Golbeck

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES

University of Maryland - College Park
Research Admin. & Advancement
University of Maryland
College Park, MD 20742 -5141

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)

ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

58352-CS-II.1

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

From July 29-30, 2010, a workshop on computing with social trust was held at the University of Maryland, College Park. Leading researchers from the diverse sub-areas relating to social trust computation came together to discuss future directions for research in the field. The goal of these discussions was to build a set of recommendations for a potential BAA in the area. This report summarizes their presentations and recommendations for areas of future work.

15. SUBJECT TERMS

Social Trust

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Jennifer Golbeck
UU	UU	UU	UU		19b. TELEPHONE NUMBER: 301-405-7185

Report Title

Summary Report for the Army Research Organization (ARO) Workshop on Social Trust Computing

ABSTRACT

From July 29-30, 2010, a workshop on computing with social trust was held at the University of Maryland, College Park. Leading researchers from the diverse sub-areas relating to social trust computation came together to discuss future directions for research in the field. The goal of these discussions was to build a set of recommendations for a potential BAA in the area. This report summarizes their presentations and recommendations for areas of future work.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Paper

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:.....	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:.....	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:	0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME PERCENT_SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Summarized presentations and recommendations for areas of future work.

Technology Transfer

Summary Report for the Army Research Organization (ARO) Workshop on Social Trust Computing

Jennifer Golbeck

September 14, 2010

From July 29-30, 2010, a workshop on computing with social trust was held at the University of Maryland, College Park. Leading researchers from the diverse sub-areas relating to social trust computation came together to discuss future directions for research in the field. The goal of these discussions was to build a set of recommendations for a potential BAA in the area. This report summarizes their presentations and recommendations for areas of future work.

Contents

1 Attendees

- Jen Golbeck (UMD)
- Ugur Kuter (UMD)
- Jim Hendler (RPI)
- John O'Donnovan (UCSB)
- Tobias Hollerer (UCSB)
- Yolanda Gil (ISI)
- Chris Hazard (NC State)
- Minindar Singh (NC State)
- Michael Reiter (UNC)
- Ronald R. Yager (Iona)
- Felix Wu (UC Davis)
- Cliff Wang & ARO researchers

2 Researcher Presentations

Slides from all presentations are included as appendices to this report. Below is a summary of the issues raised by the participants.

2.1 Cliff Wang

Discussion Goals

1. How is this topic multidisciplinary? Be specific.
2. Why is this topic basic research and not technology development?
3. What scientific understanding is lacking in this topic area?
4. Why has that understanding not been obtainable? What are the difficulties?
5. What recent scientific breakthroughs or accomplishments that would now enable us to do so?

6. Why more research is needed to achieve that understanding?

7. By what criteria will the research be judged to be successful?

3, 4, 5 are particularly important. What are the fundamental things we need to achieve?

Information used to come from sensors, but now it comes from a combination of sensors and humans, and there is contradictory information. More and more information is coming in from social interactions. How do you make a decision? What is a scientific method for doing that?

Can we trace information back to the origin? (This raises the provenance question)

2.2 Ron Yager

Title: Fuzzy and Soft Technologies for Trust

Many aspects of social trust computing require the processing of information that is imprecise, uncertain and subjective which can benefit from the use fuzzy and other soft computing technologies. The focus of our talk will be on the role that fuzzy sets and other soft granular technologies can play in furthering trust modeling and research. We shall describe an extension of social network modeling that uses fuzzy sets to build a bridge between the types of linguistic terms which humans use to communicate, reason, and understand and the formal representation of the social networks. This allows for the human linguistic articulation of trust related concepts to be translated into computational manageable mathematical objects. An important task in trust computing is aggregation, this occurs both in the determination of reputation as well as in path based process of determining trust via social network connections. We describe a number of semantically based aggregation tools developed in the soft computing framework and discuss the applicability to trust computing. These will allow one to process linguistically expressed measures of trust in addition to numeric ones. Closely related to this is the technology of linguistic summaries that allows one to formulate and manipulate statements such as "In most cases John is a highly trustworthy partner." We shall briefly discuss the paradigm of participatory learning which can provide a framework for context based trust determination.

2.3 Mike Reiter

Title: Leveraging Community for Usable Security

The proliferation of always-connected mobile devices permits one's community to be reached at nearly any moment. We will present ideas for using this pervasive community presence to manage security policies, to exercise one's authority, and to protect the security of one's own information. Part of this talk will draw from our experience with Grey, a system we have deployed on two university campuses to enable users to leverage their mobile devices for exercising and delegating authority to physical and virtual resources.

2.4 Felix Wu

Title: "SCOSK: Social-Centric Operating System Kernel"

In this short talk, I will discuss the possibility of pushing social informatics into the Operating System kernel, mainly for the purpose of cyber security. We will first go over a few cyber events/attacks as motivation examples. Then, we will discuss the benefits of managing social resources, as a first-class citizen, inside an OS kernel. Finally, if time allows, we will go over a few technical/research challenges in realizing SCOSK.

2.5 Munindar Singh

Title: Trust Middleware as a Basis for Social Computing

Recent research into trust has yielded approaches that consider evidence, incentives, cognitive states, and social relationships.

But the key aspect of trust that underlies the above approaches is itself largely ignored. Specifically, trust reflects a dependence of one agent on another for a purpose. The mutual dependence of agents and their successes or failures pertaining to it may be reflected in social relationships, expressed cognitively, motivated by incentives, or recorded in evidence. But the representation and reasoning about dependence is a central concern that merits serious study.

Based on the above, we motivate a trust middleware as a central logical component for social computing. The middleware would assist agents by bookkeeping their trust relationships. It would help realize application-specific architectures that support diverse social computing applications, such as those involving personal, communal, organizational, and contractual relationships.

2.6 Ugur Kuter

Every day, people and organizations exchange information and make decisions through networked entities. Sometimes the nodes of these networks

may be other people or organizations; sometimes, they are communication hubs, Internet servers, or software agents. In any case, the activities, goals, and plans of any entity depend on the trust and credibility of the networked entities and the information they spread across the network. Such notions of trust and credibility also affect how certain events happen in the world, which are not under control of the decision-making agent. Thus, it is important to understand how agents or groups of agents use resources in the world, how they produce information based on those resources, and how those resources determine their plans, activities, and goals over time. To date, social dynamics in networks have been largely neglected in systems that help to make decisions under spatial and temporal constraints. I'll describe some background and the works we're doing currently at UMd in this specific research topic. I'll then summarize several future directions.

2.7 Chris Hazard (presented by Minindar Singh)

Title: Trust and Incentives: Dynamics and Strategy

Abstract: In multiagent interactions, such as e-commerce and peer-to-peer file sharing, being able to accurately assess the trustworthiness of other agents is important for agents to protect themselves from losing utility. We focus on an agents' discount factors (time preference of utility) as a direct measure of the agent's trustworthiness in a number of settings. We prove that an agent's discount factor, when in context of the agent's valuations and capabilities, is isomorphic to trustworthiness for a set of reasonably general assumptions and definitions. Further, despite the large body of work in reputation and trust in dynamic multiagent environments, no metrics exist to directly and quantitatively evaluate and compare reputation systems. We present a common conceptual interface for reputation systems and a set of four measurable desiderata that are broadly applicable across multiple domains, inspired by dynamical systems theory. We discuss the implications, strengths, and limitations of our desiderata. Our discount factor trustworthiness model performs well across the desiderata when measured against other established reputation models from the literature. We briefly conclude with some preliminary work extending our dynamics and strategy perspective into the dimensions of signaling and sanctioning as a way of designing and assessing trust systems.

2.8 John O'Donovan

Based at the Four Eyes Lab, Dept. of Comp Science, University of California Santa Barbara, the WiGis project centers around visualization of large-scale, highly interactive graphs in a user's web browser. Our software is delivered natively in your web browser and does not require any plug-ins or add-ons. Our method produces clean, smooth animation in a browser through asynchronous data transfer (AJAX), and access to rich server side resources without the need for technologies such as Flash, Java Applets, Flex or Silverlight. We believe that our new techniques have broad reaching potential across the web. Below is a windowed version of the system.

The WiGis framework has been applied as a solution in a diverse set of scientific, social, computational and other systems that require user interaction with networked data. Please see the left menu for some examples and live demos, or click the button below to examine the core framework on some static data.

2.9 CTA discussion

The trust research taking place in the Army's centers for network science is a vast undertaking. Nearly all of the attendees are affiliated with the CTA effort. A discussion of the research areas covered in the different centers and cross-cutting research initiatives identified topics being covered and those not addressed by the centers. The following research questions are currently not part of the work taking place under CTA:

1. Non-network centric work - Since the CTA is focused on network science, non-network related trust issues are outside that domain. Of particular interest to the workshop attendees are behavioral and linguistic trust models. Behavioral models in particular are part of a rich research history that has new applications to the wealth of data from online social interactions. A sub-topic discussed in this area is the motivation to trust. What causes people to trust and what do they see as its benefits?

2. Dynamics of trust - This topic addresses how do trust values change, if forgiveness is considered, the changing justifications for trust levels, etc. Trust dynamics are of interest to many groups but are not well studied, so this is a question with lots of potential.

3. Distrust / negative trust - There have been a few studies on distrust or negative trust values, but it is problematic to use in many models.

However, like trust dynamics, it is an issue of great interest to the military and to trust researchers.

4. Trust risk model - A security model, where trust is bad and security can be compromised is quite different from a social model where trust is good because it helps with decision making and where risk is not as clearly defined. Looking at the overlap between these tracks of research is an open space and may lead to some interesting integrations of theories.

5. Use of game theory - Trust and game theory fit well together and there is extensive literature on the topic. It is still an active area of work where new results are being produced, and this is not a topic of study in the CTA.

3 Ideas for Future Trust Research

After the presentations and discussions presented above, we discussed twelve areas that emerged as spaces for future work. Each topic discussion was led by a workshop attendee whose name is indicated parenthetically next to the topic. In the section below, each topic is described with possible research areas. Where appropriate, extra attention is given to the questions raised above by Cliff Wang regarding the research that's missing, why we can do it, and what the impact will be.

3.1 Dynamics of trust

(Jinghe Cho, Munindar)

How does trust change over time? What are the roles forgiveness, forgetfulness, regretfulness? How quickly do a person's actions propagate and affect trust that others have in them? How does that affect others in the network?

There are countless questions related to the dynamics of trust, and it is an area where there was a lot of enthusiasm at the workshop.

On a system level, there are issues of how the system impacts trust among users, how trust is represented, and how trust values are updated. Algorithmically, there are questions about how trust dynamics are represented in algorithms and how changes propagate through the network.

Finally, there are behavioral issues related to how dynamics affect and are affected by social structure and social groups.

This area of trust dynamics touches on many of the open areas not covered by CTA, and is one that the workshop attendees saw as a space that many of us could do work.

3.2 Trust/risk model

(Kevin)

Trust is a relationship that people use to mitigate risk. However, the study of risk itself is of interest to better understand trust. Furthermore, as mentioned above, risk plays an important part in security discussions, and this is a space where social and security oriented risk discussions can come together. Research topics in this space include the impact on decision making, uncertainty management, time sensitivity of decisions with respect to trust, and methods for minimizing risk and optimizing decision making.

3.3 Linguistics

(Presented by Ron, Chris)

The use of language for understanding trust has potential, particularly in the area of ontologies, topic modeling, and applications of LDA.

We discussed using computational linguistics as a mechanism for weighting trust, and for extracting trust from statements.

Linguistics also have potential for revealing why people have trust in others. It provides insight into personal interpretation of people and events that lead to trust. The technologies to pursue this work now exist, but they have not been extensively applied to study trust.

3.4 Human Behavior and Trust

(Ugur)

Studying the relationship between human behavior and trust is a wide open space for research. Looking at observed behavior - do actions indicate trusting or untrusting behavior - can be studied in many ways. Computational social science, game theory, using controlled experiments with users, and observing their behavior are all methods for doing research in this space.

Behavior and trust will also touch on questions of background and culture, which is of particular interest to military applications. It will also cover questions of impersonation, deception, and lying.

3.5 Trust aggregation and inference

(Ron, Ugur)

Trust inference is one of the more widely studied areas of trust. There are a number of algorithms published and implemented in systems. However, there are many questions left to be studied in this area. To date, there is no research on how to combine trust values from different sources, how to allow user control over algorithmic functions (e.g. should people have personalized trust scales, could and should trust aggregation functions vary between users).

Evaluation of these algorithms is also not standardized, partially because of the lack of public trust datasets. If a BAA is to encourage research in this space, evaluation should be encouraged and datasets should be provided to allow direct comparison of performance between algorithms.

3.6 Engineering trust systems

(Munindar)

While most of our discussions focused around theoretical and scientific research questions related to trust, there was limited interest in the engineering of systems. Munindar, who lead this discussion, raised issues of a trust middleware, centralized trust systems, and the ability to use trust across systems. This is a very different topic than the rest suggested by the workshop attendees, but could be an interesting area for research if a more engineering / system building approach is desirable.

3.7 Properties of trust systems

(Felix Wu)

Related to the engineering of trust systems, but a topic that received a more enthusiastic response, is that of properties of trust systems. This includes understanding global properties, like robustness (including anti-manipulation / stabilization), achieving consensus, preserving diversity and protecting minorities, confidentiality, accountability, and incentives for participation.

Robustness brings us to issues of attack resistance. In trust systems, that will include somewhat traditional attack strategies, like collusions, but also questions of how social influence or reputation can be exploited to take advantage of others in the network. On the flip side, study of offensive strategies for disrupting trust networks falls within this area. This is useful for learning how to protect a network and for harming adversaries.

3.8 Leveraging trust in systems

(Mike)

One of the main reasons for studying and understanding trust is to use it in applications. While there has been some study of the effectiveness of trust in applications, there is a lot of exciting work to be done in this area. Application spaces where trust is used to improve other systems include recommender systems, trust in automation, and information filtering and personalization. Trust can also be reapplied to social systems; strengthening interactions, recommending relationships, and investigating the connection between trusting individuals and groups are all important research questions with many open problems.

3.9 Communication of trust

(John)

Once trust is used in applications, it may appear in back end systems like recommender systems or on front-end interface elements that indicate how trustworthy a person or their content is. For users to take advantage of the trust, there are questions of how (if at all) trust should be communicated to the end user *and* how users provide trust information to the system.

Research on communication of trust includes studying what is the right amount of trust to communicate to users (black box models vs. transparent models) and how entities in the system exchange information with one another. Research in this space will also allow comparison and performance differences between various trust models more easily.

3.10 Distrust

(Jen)

Distrust is an understudied question in this area. During the workshop, we raised a number of research issues including propagation and representation of distrust. However, while the workshop attendees recognized this as an important topic, there was little enthusiasm for this as an emerging area. Part of the reason it is not well studied now is that the unaddressed issues are quite complex and difficult to work with properly. Some argued that it may be impossible to use distrust in more than a very simple way. We felt it was unlikely that a BAA would lead to much progress on this topic.

3.11 Context of trust -

(Ron)

Context is an issue that is constantly raised in discussions of trust. It is certainly an important issue, as both the meaning of trust and its treatment in algorithms and models will vary as context changes. However, while we all recognized that it was important, there were no real research questions that emerged nor enthusiasm about this as a topic for the BAA.

THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Leveraging Community for Usable Security

Michael Reiter

Lawrence M. Slifkin Distinguished Professor
Department of Computer Science
University of North Carolina at Chapel Hill

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

High-Level Message

- With proliferation of always-connected mobile devices, your community is nearly always accessible to you
- Can we make use of them to help with our own security?

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Background: Grey
[w/ Bauer et al. 2004-present]

- Grey is an implementation of credentials-based access control
- Utilizes smartphones as the means by which users exercise/delegate authority
- Grounded in formal logic to permit rigorous proof checking and assurance arguments
  - Also central to supporting *reactive delegation*

An Example Scenario

UNC can access online journals

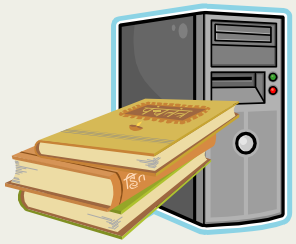

ACM

- All members of the car
  can access ACM's onlin
- Alice is a member of th
  community

UNC Admin

- Alice is writing the related work section of a paper and needs access to an online journal

Alice

An Example Scenario

ACM

UNC Admin

Alice

An Example Scenario

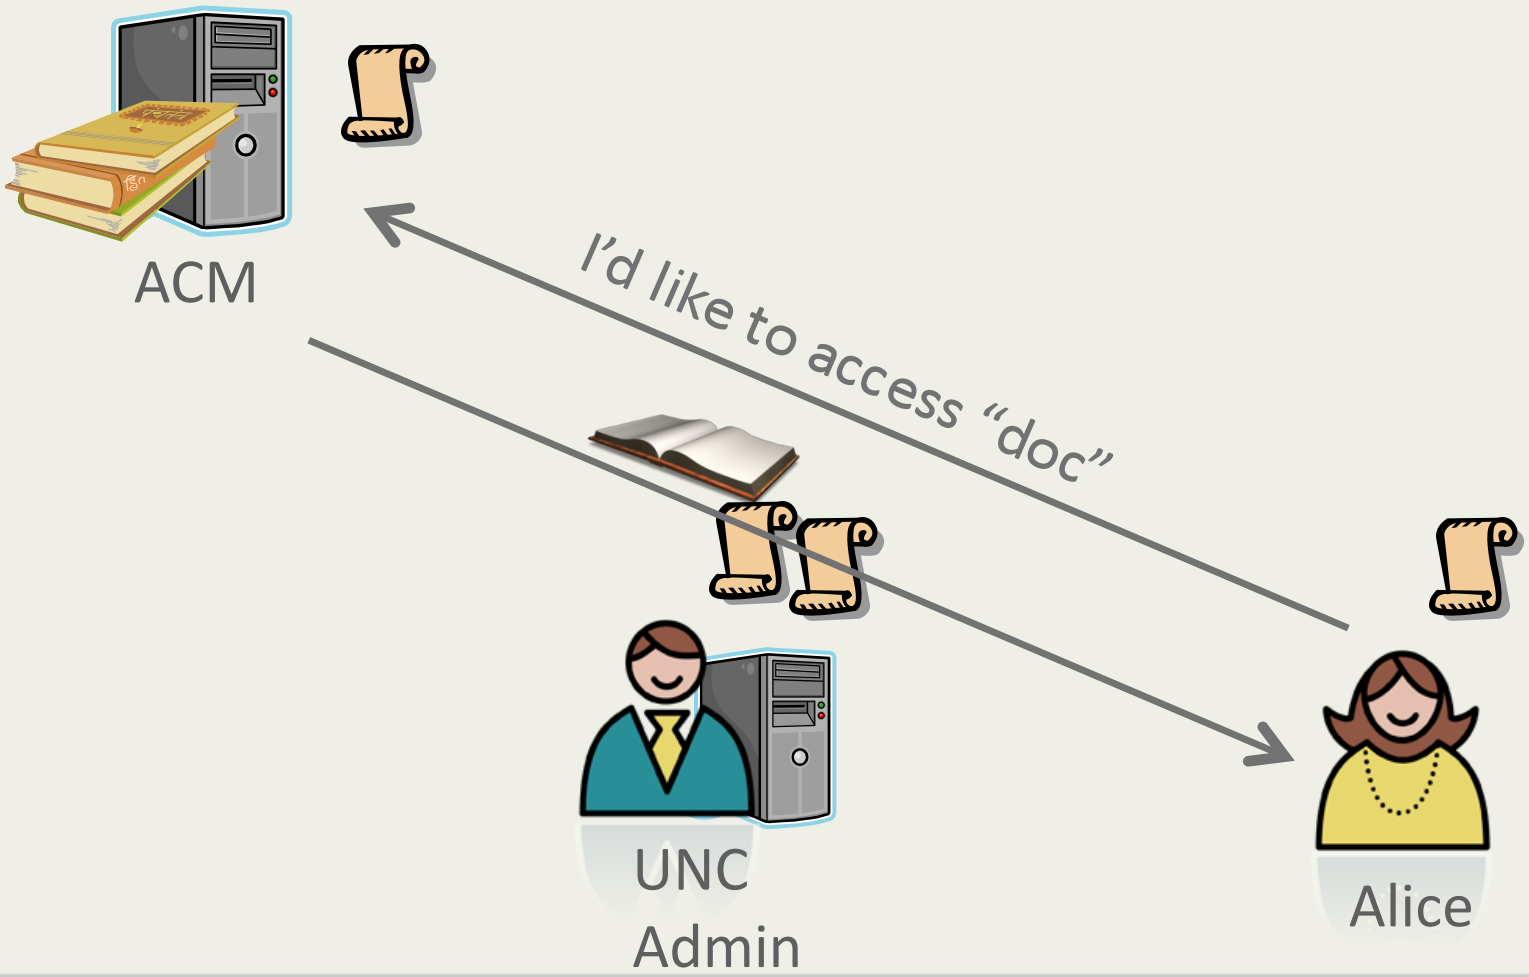

Reactive Delegation
[w/ Bauer, Garriss 2005, 2007]

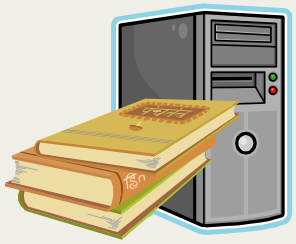

ACM

Where can I get credentials to finish the proof?

I'd like to access "doc"

What credentials should I create?

UNC Admin

Help!

Alice

Grey Deployments

- Two deployments for physical access control
  - CMU's Collaborative Innovation Center
  - UNC's Fred Brooks Building

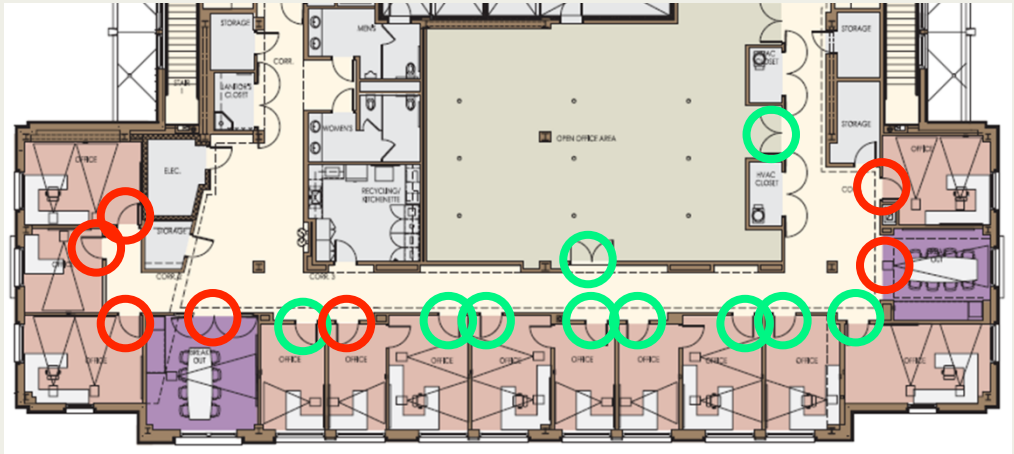

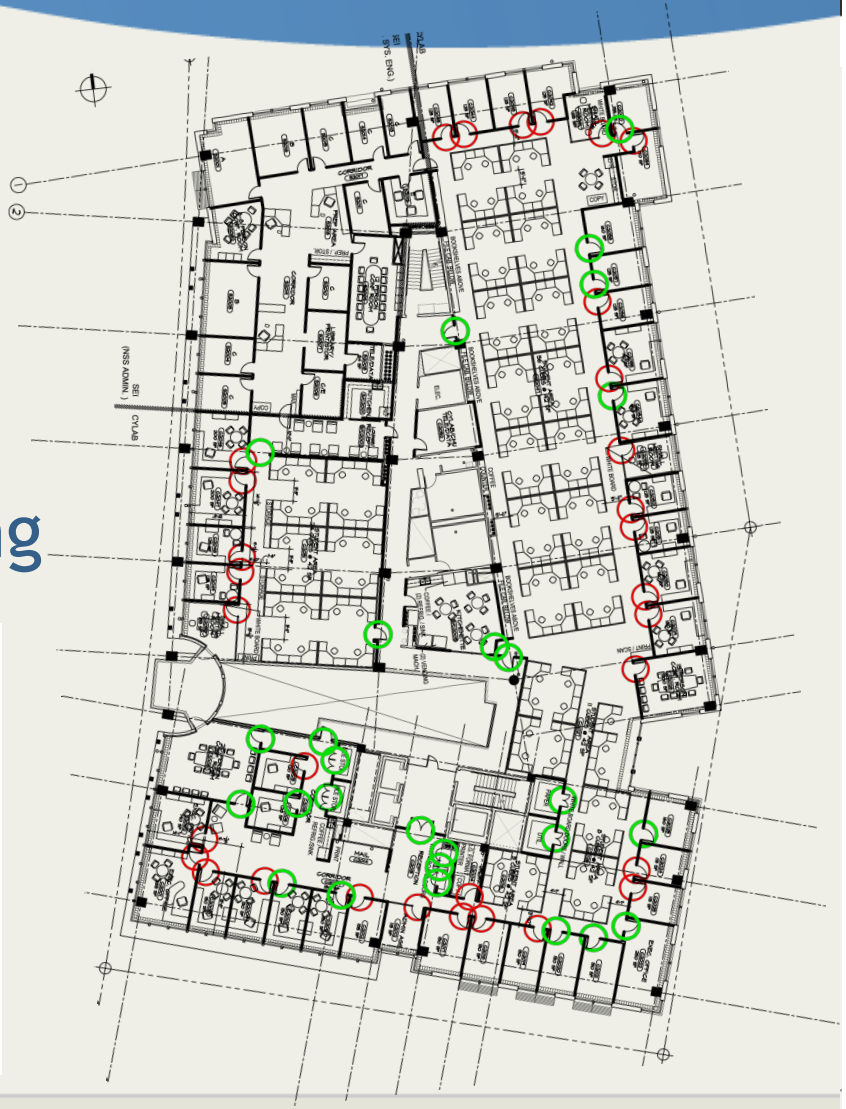

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Benefits of Reactive Delegation
[w/ Bauer, Cranor, Reeder, Vaniea 2007-8]

- Grey policies match user intentions better than those implemented by physical keys
- Inflexibility of physical keys caused people to implement "workarounds"
  - E.g., a "key drawer" where keys were stored (unlocked!) for those who knew where to look
- On-demand (reactive) delegation permitted a "least privilege" posture
  - Without undue risk of unavailability

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Leveraging Community

- Reactive delegation is one instance of leveraging community to manage security
- Key question: Are there other applications in which community management can improve security with little overhead or inconvenience?
  - Sure, reputation systems …
  - … But how about something more personal?

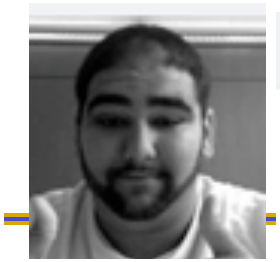
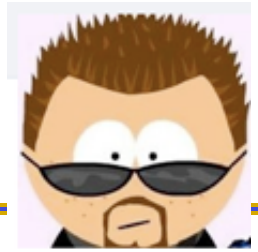


UC DAVIS
UNIVERSITY OF CALIFORNIA

DSL

Davis Social Links
SCOSK: Social-Centric OS Kernel

S. Felix Wu
Computer Science Department
University of California, Davis
wu@cs.ucdavis.edu
http://www.facebook.com/sfelixwu
http://dsl.cs.ucdavis.edu

SCOSK

UCDAVIS
UNIVERSITY OF CALIFORNIA

- Social Informatics
  - Understand/Utilize ==> Manipulation
    - Social Network has a dynamic perspective
    - One single giant SN versus per-application
  - Social Informatics inside the OS kernel?

- Application Context
  - Decision Making, Policy Enforcement, Community Development
  - Collaborative Tasks, Information Exchange

We are trying to make the information of social context and relationship explicit!

UC DAVIS
UNIVERSITY OF CALIFORNIA

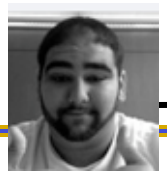

SMTP

facebook Home Profile Friends Inbox 575 S. Felix Wu Settings Logout

Send Justin a Message

Information

Relationship Status:
Single

Friends

26 friends See All


Cristy Alstrand


Casey Hentzen


Leslie Almonte


Daniel Ramon


Amanda Martinez


Sara Royer

Justin Howell Awaiting friend confirmation

Wall Info

Justin only shares some of his profile information with everyone.
If you know Justin, send him a message.

Basic Information

Sex: Male
Relationship Status: Single

Personal Information

About Me: Just taking life one beer at a time. Or twelve.

Education and Work

Colleges: UC Davis '08
Computer Science
Solano CC '05
7 A.S. Degrees, bitches!
High School: Fairfield High School '98
Employer: Solano Community COllege
Position: Telecommunications/Network Technician
Time Period: August 2002 – Present
Location: Fairfield, CA
Description: Yeah I run this place

Pages See All (7)


Sriracha (Rooster Sauce)

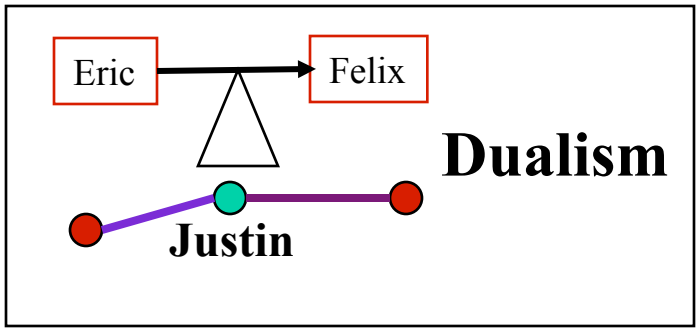

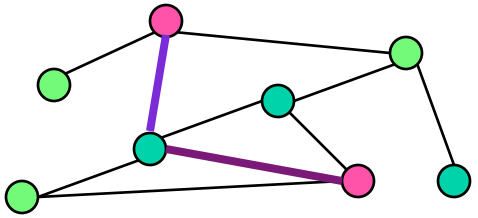

0.65

0.73

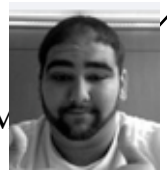

SquirrelMail version 1.4.19
By the SquirrelMail Project Team

Davis Social Links Login

Full Email: sfelixwu@gmail.com

Password: ••••••••••

Login

Log in to: Gmail

Remember login server

Instructions:

1. Join CyrusDSL
2. Register your email with DSL (NOTE: You must type in your full email address, i.e. "yourname@gmail.com".)
3. Log in above with your full email address and it's password.

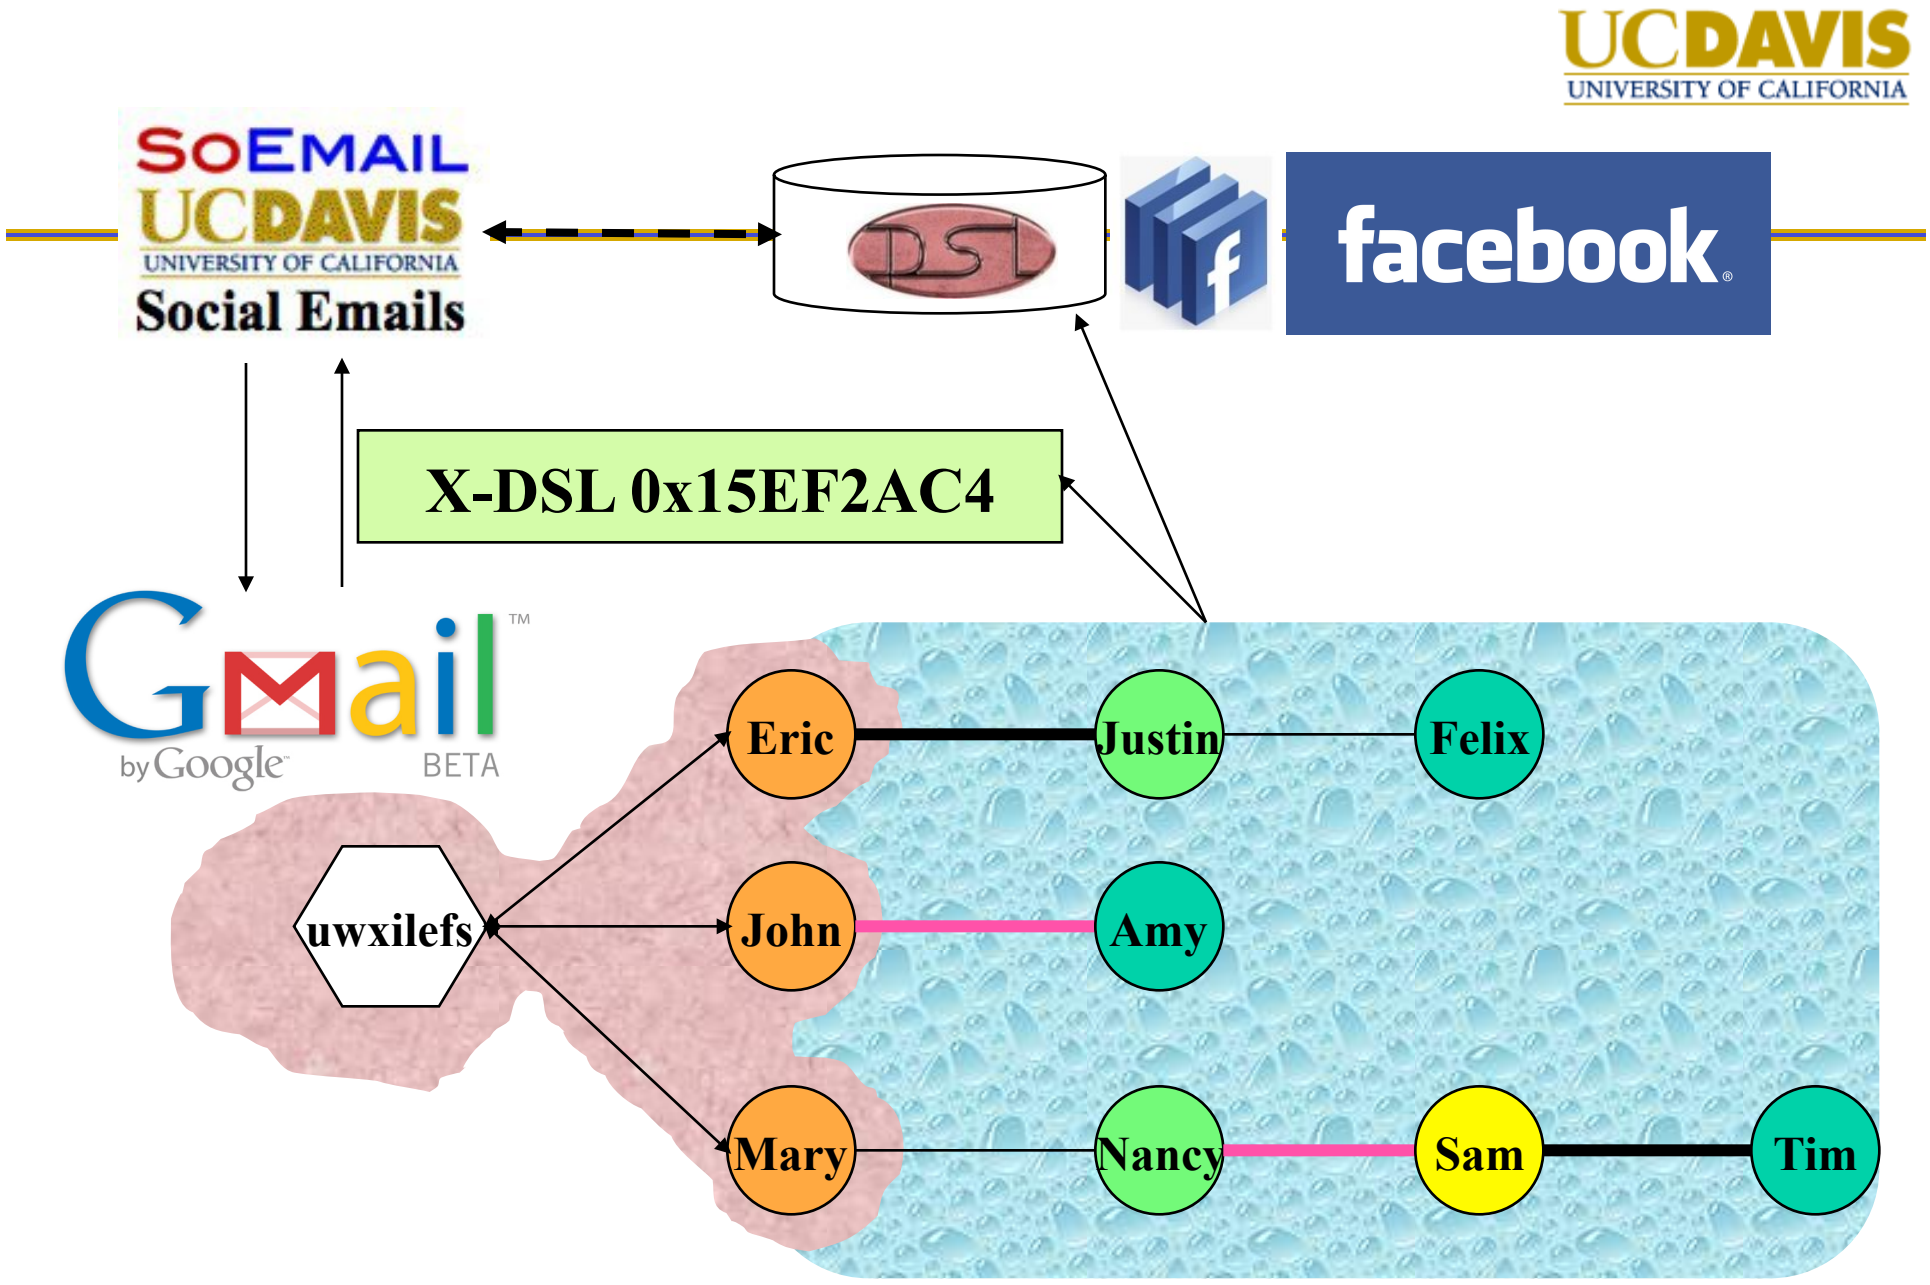
SoEMAIL
UCDAVIS
UNIVERSITY OF CALIFORNIA
Social Emails
DSL
facebook
X-DSL 0x15EF2AC4
Gmail
by Google
BETA
uwxilefs
Eric
Justin
Felix
John
Amy
Mary
Nancy
Sam
Tim

OS Kernel

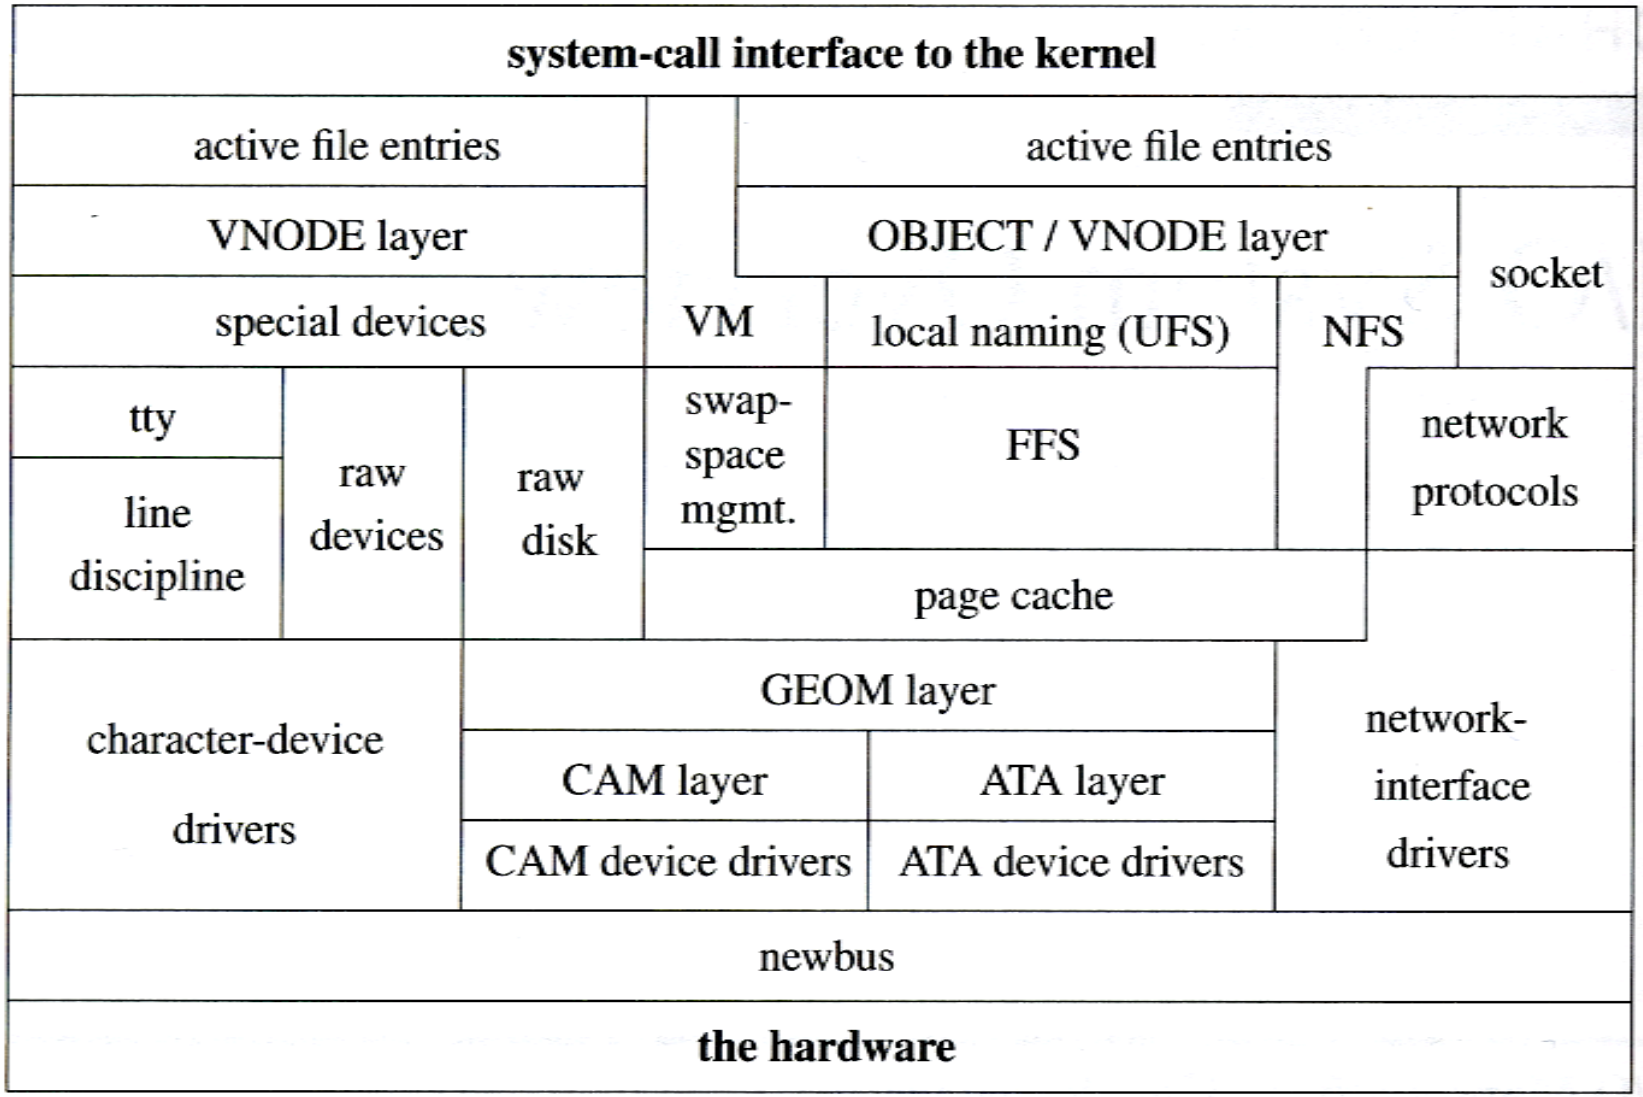
system-call interface to the kernel
active file entries
VNODE layer
special devices
active file entries
OBJECT / VNODE layer
socket
VM
local naming (UFS)
NFS
tty
line discipline
raw devices
raw disk
swap-space mgmt.
FFS
network protocols
page cache
character-device drivers
GEOM layer
CAM layer
ATA layer
CAM device drivers
ATA device drivers
network-interface drivers
newbus
the hardware

OS Kernel

- It's all about resources
  - CPU Cycles, Memory, File System, I/O Bandwidth…
  - Application Processes!
- Allocation and Abstraction
  - "Fairness/Efficiency"
  - "Provenance"

Social-Centric OS Kernel

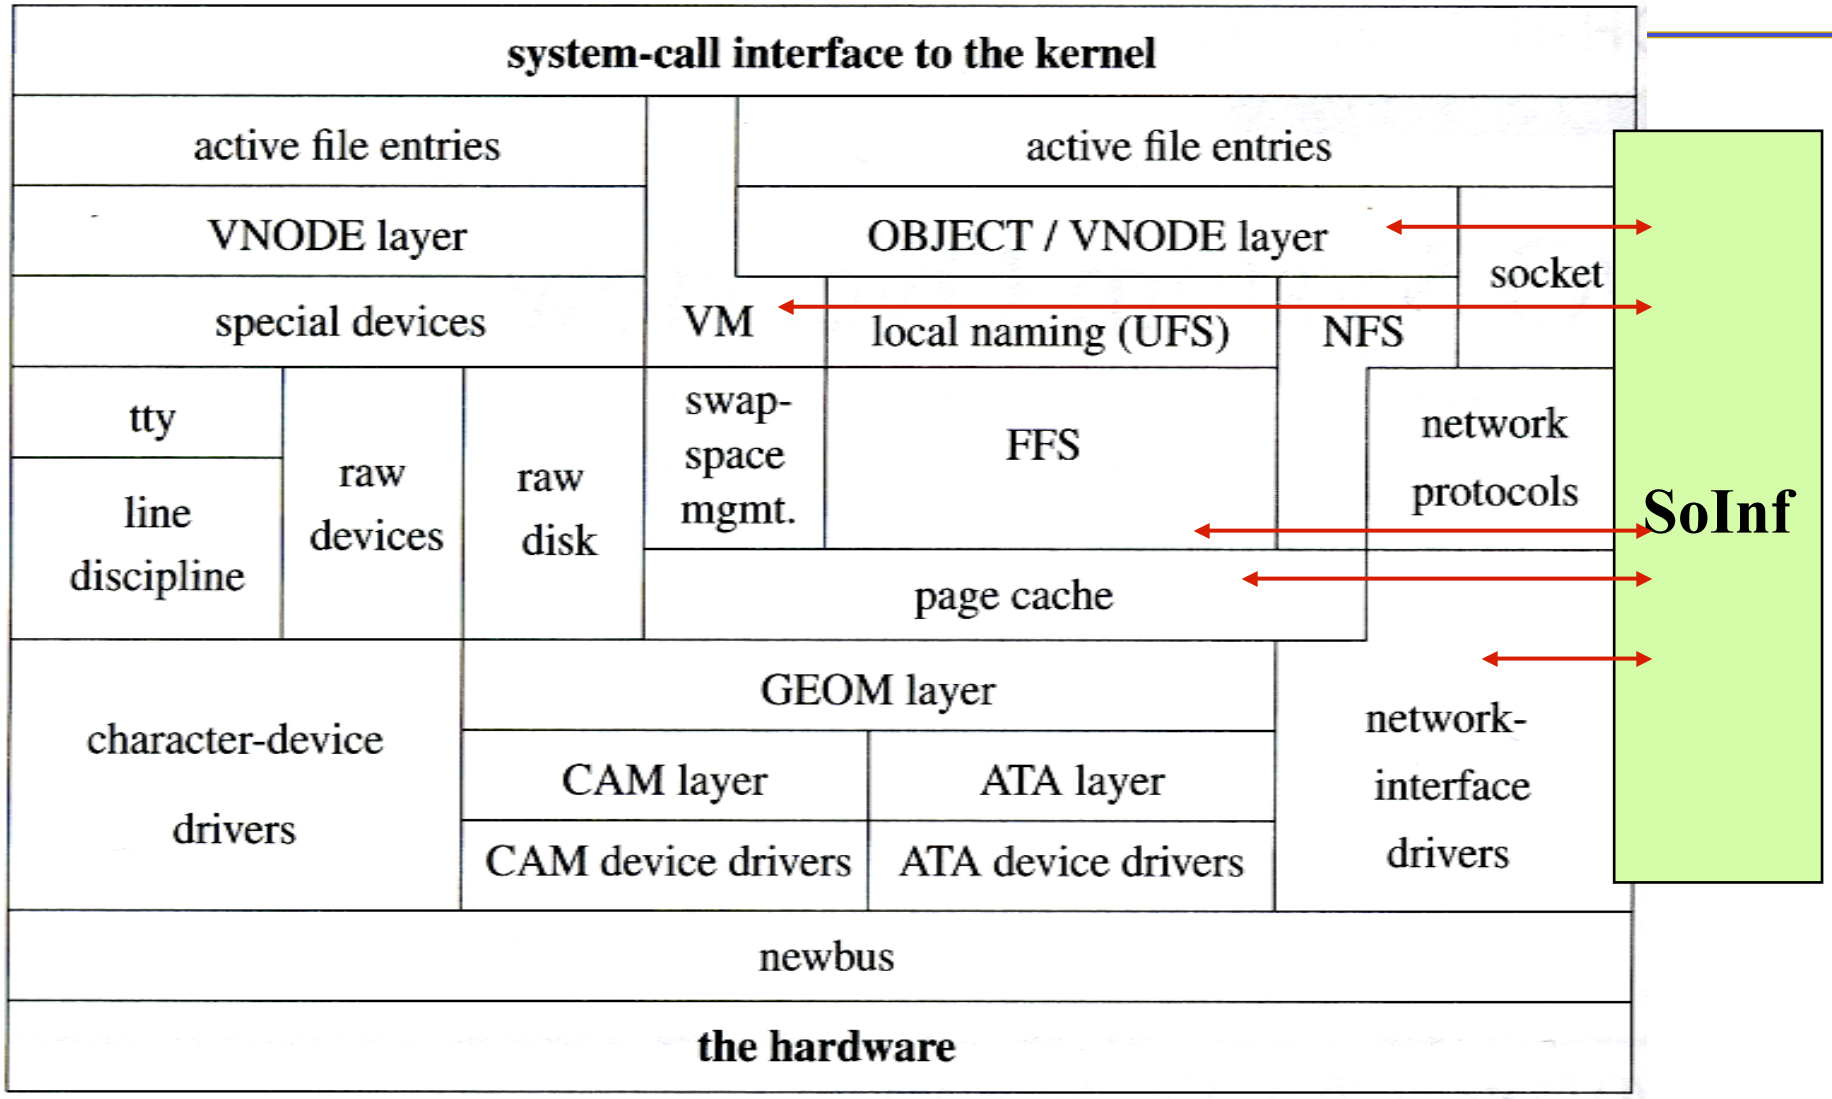

OS Kernel

- It's all about resources
  - CPU Cycles, Memory, File System, I/O Bandwidth…
  - Application Processes!
- Allocation and Abstraction
  - "Fairness/Efficiency"
  - "Provenance"
- How do the SCOSK protect the core value of the SI against aggressive applications?

Farmville and Facebook

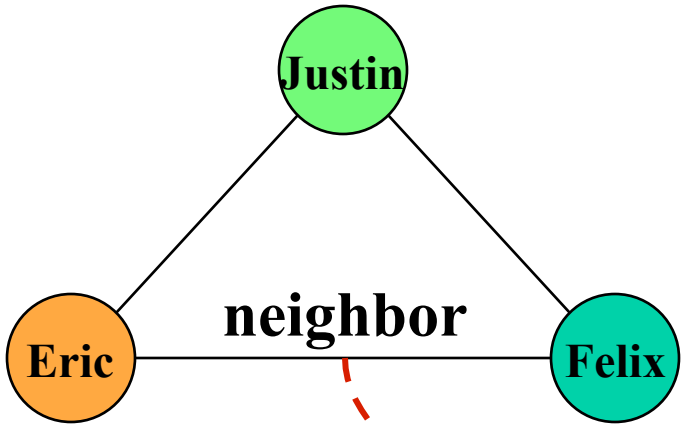

FB friends: 790+
FV neighbors: 30+

The "Add me please" push!!!

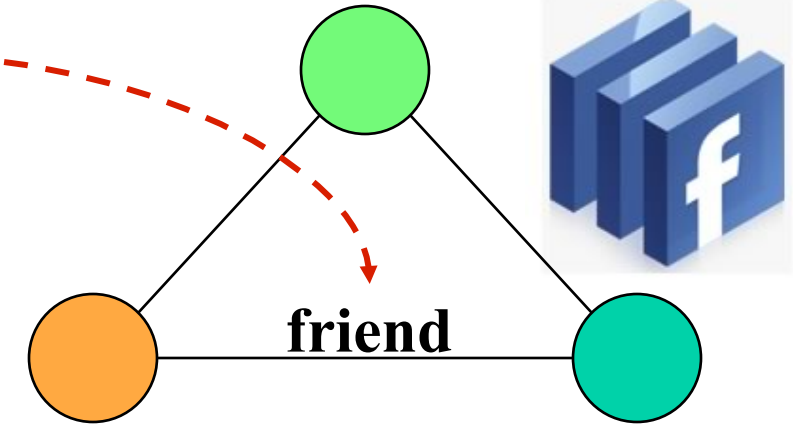

A Couple Issues

- Careless in adding friends regardless …
  - Incentive model at the "Facebook" level to balance such a behavior
- Certain Applications are pushing users to accept new friendships unconditionally.
  - But, how do we control the quality of the friendships, at least, in the context of a particular application?

A range of possibilities…

- Bit Torrent Peer selection
  - Currently no "Social" Relationship involved
    ==> maybe "Reversed OSN"
- Farmville Neighbors
  - Must leverage Facebook friendship

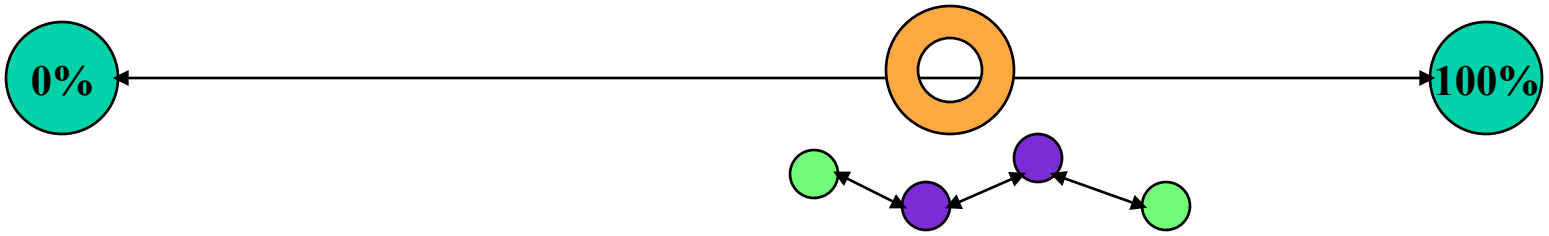

What is the appropriate Social Informatics for application X?

Farmville and Facebook

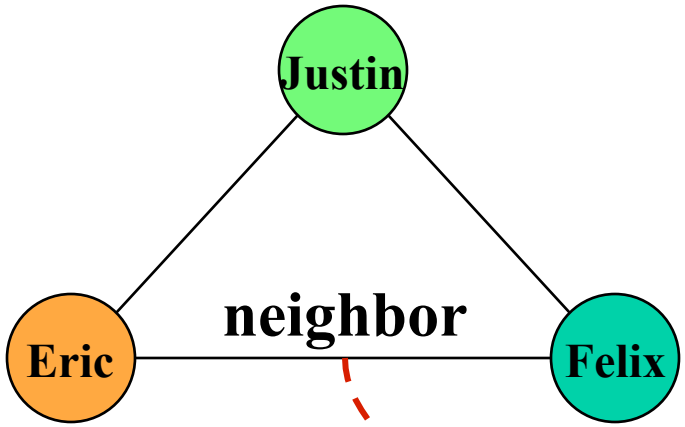

FB friends: 790+
FV neighbors: 30+

Is this necessary?

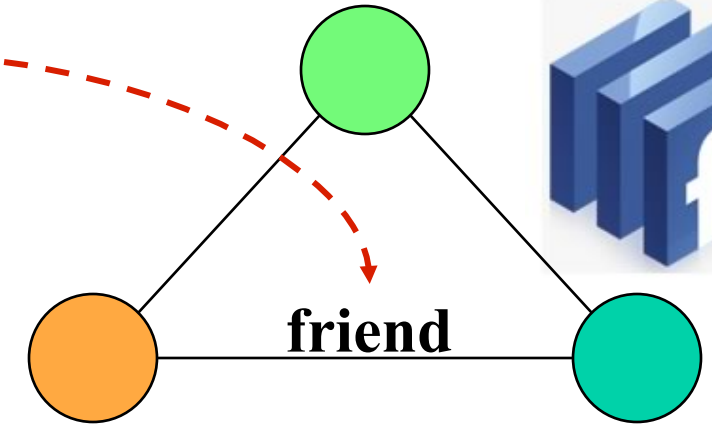

Let's review…

- How social informatics is being accessed under Facebook?

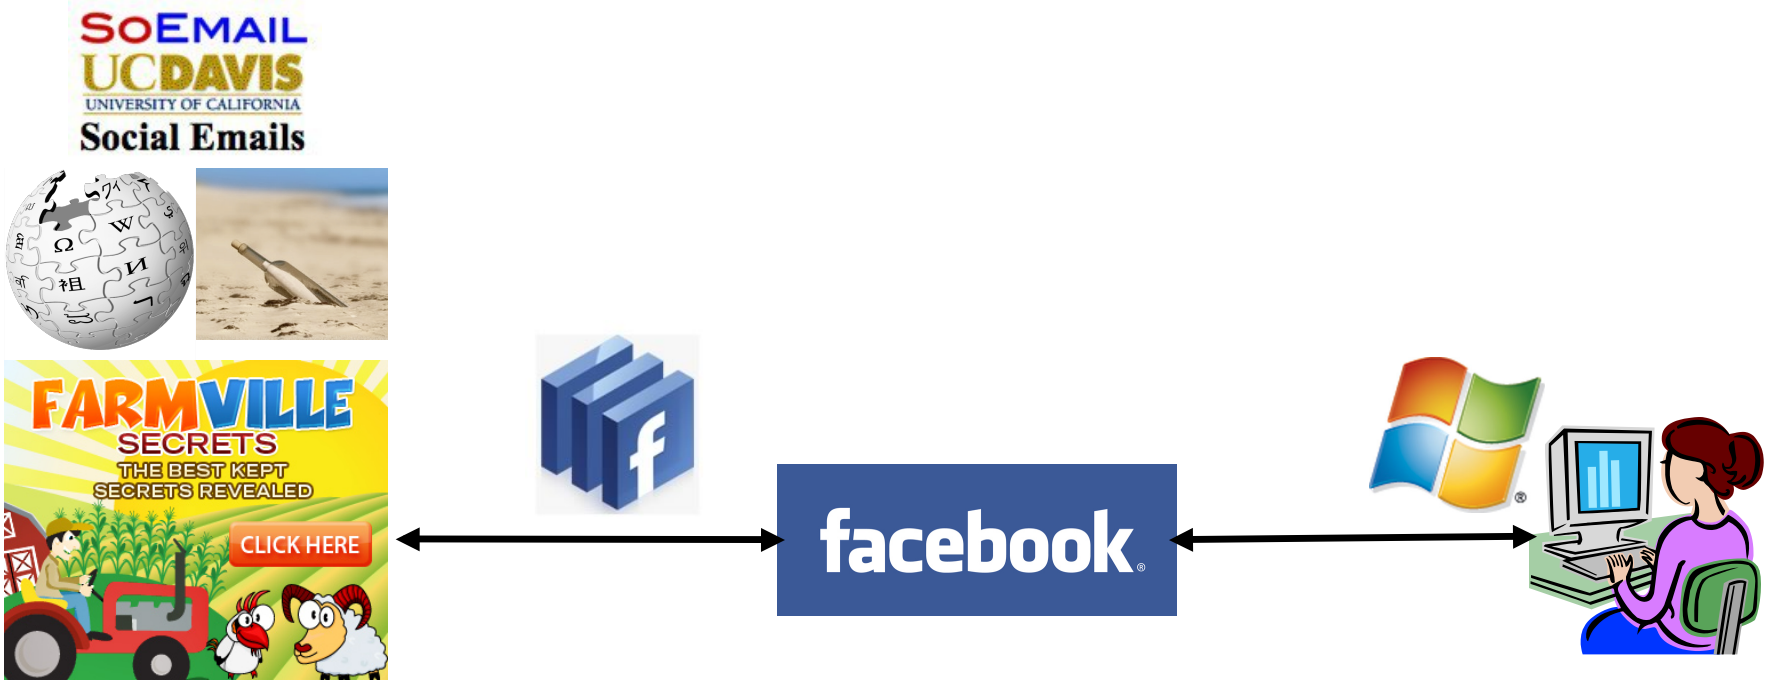
SoEMAIL
UCDAVIS
UNIVERSITY OF CALIFORNIA
Social Emails
FARMVILLE
SECRETS
THE BEST KEPT
SECRETS REVEALED
CLICK HERE
facebook

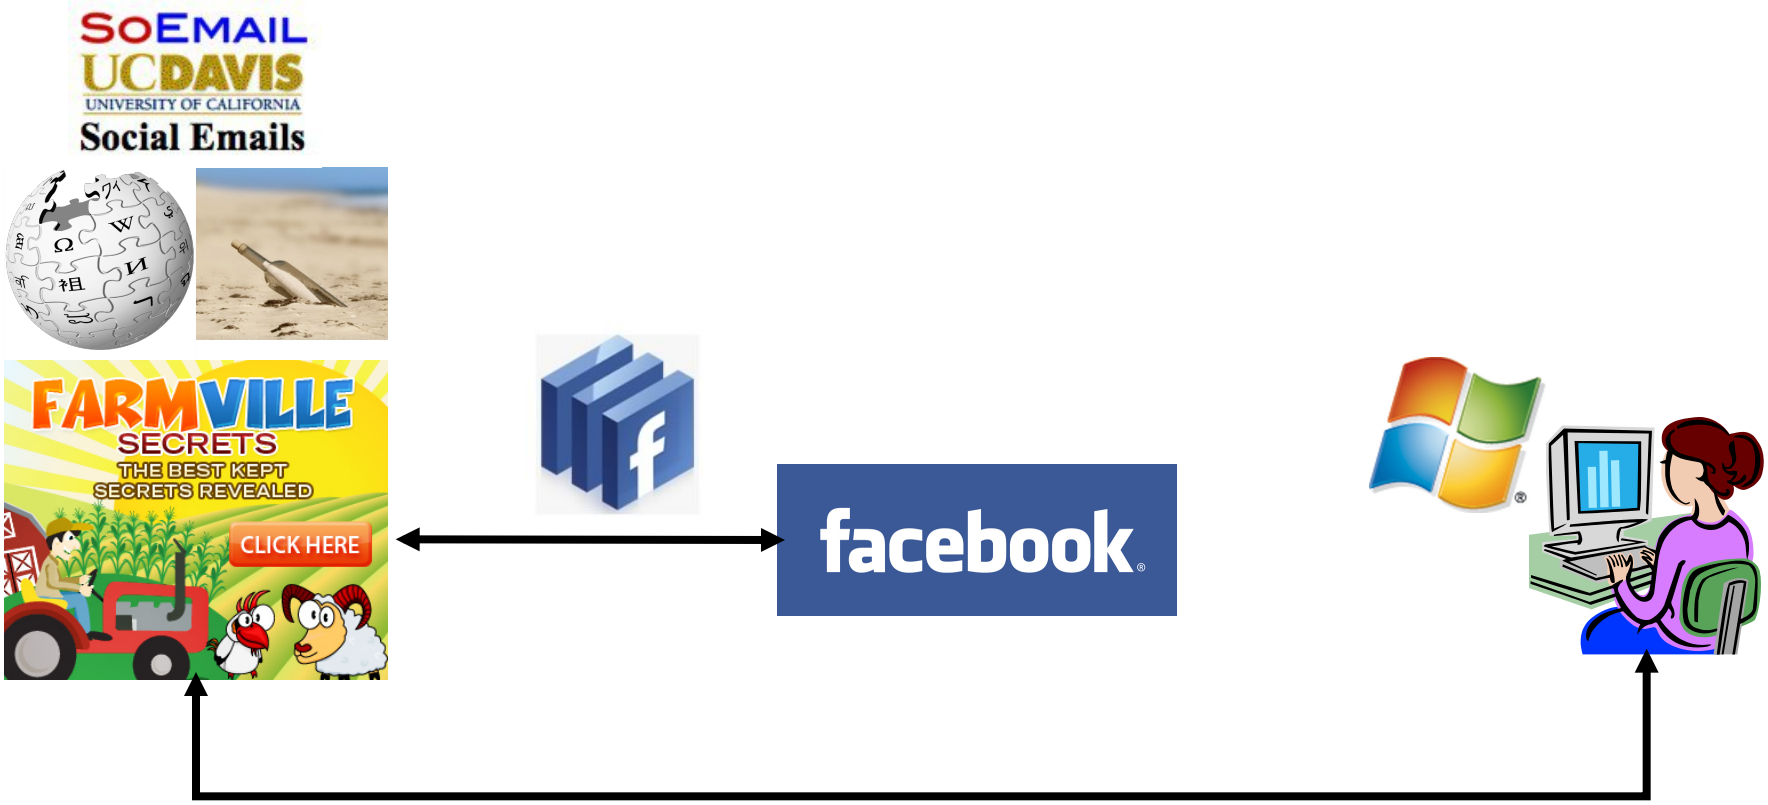
SOEMAIL
UCDAVIS
UNIVERSITY OF CALIFORNIA
Social Emails
FARMVILLE
SECRETS
THE BEST KEPT
SECRETS REVEALED
CLICK HERE
facebook

FB-connect

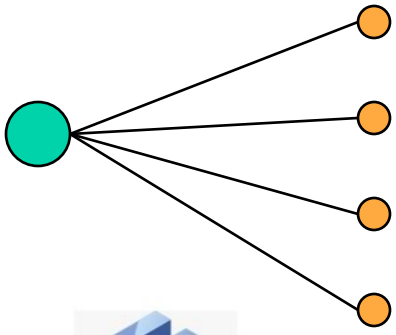

Problems:
-aggressive pushing
-Hard to control from the user's perspective

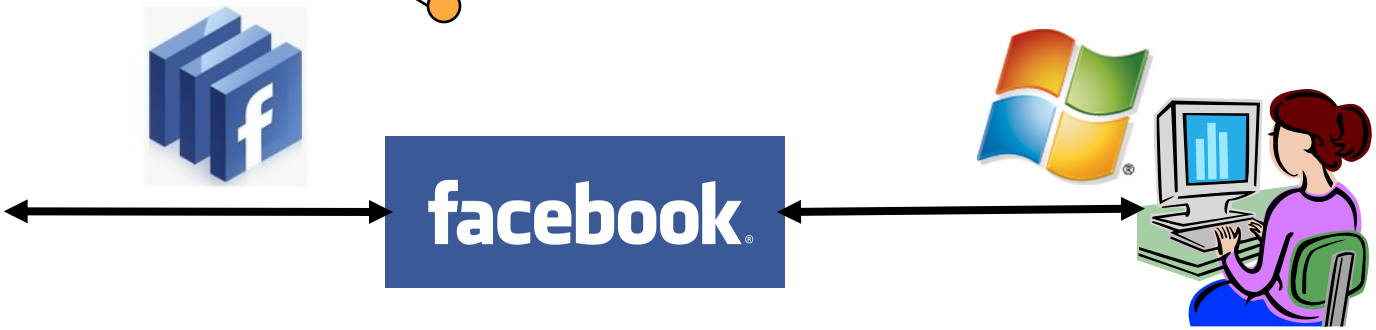


```
friends.get
friends.getappusers
```

UCDAVIS
UNIVERSITY OF CALIFORNIA

FAITH

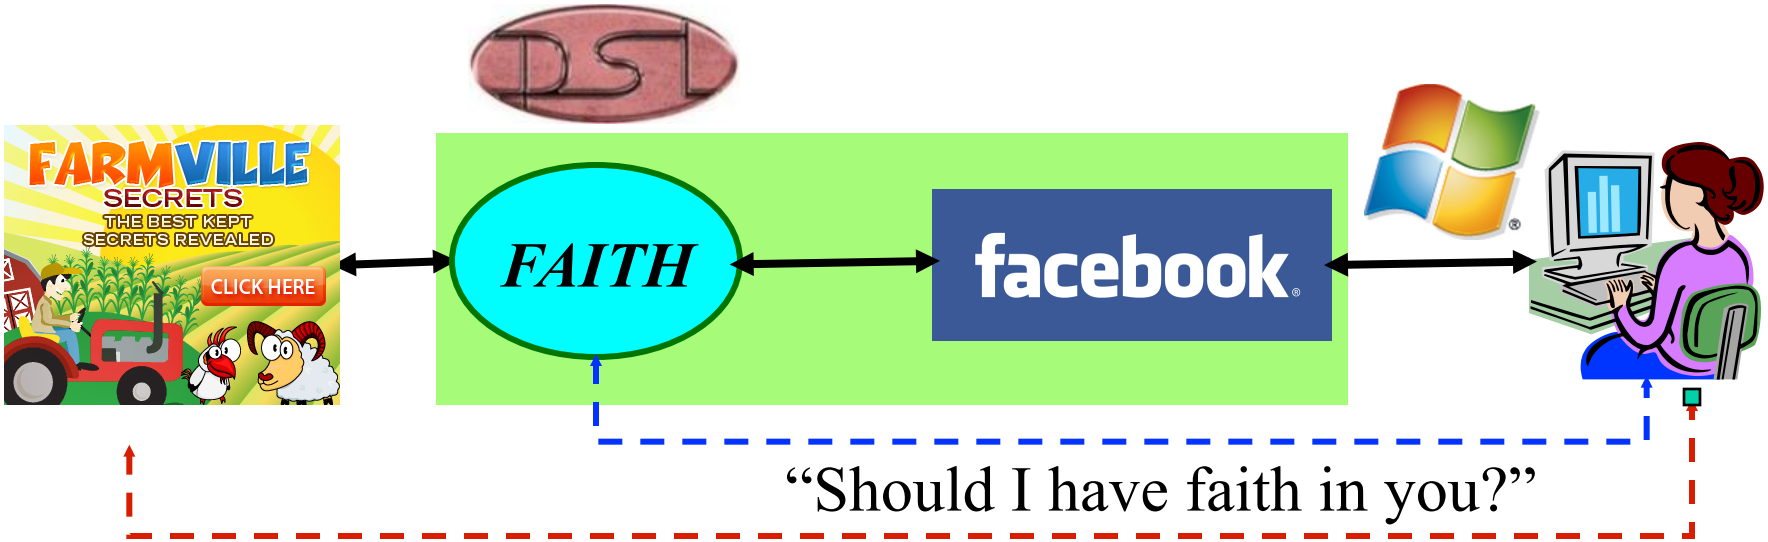

FAITH (**F**acebook **A**pplication **I**dentifier **T**ranslator & **H**ypervisor)
like *NAT (Network Address Translation)*

Farmville and Facebook

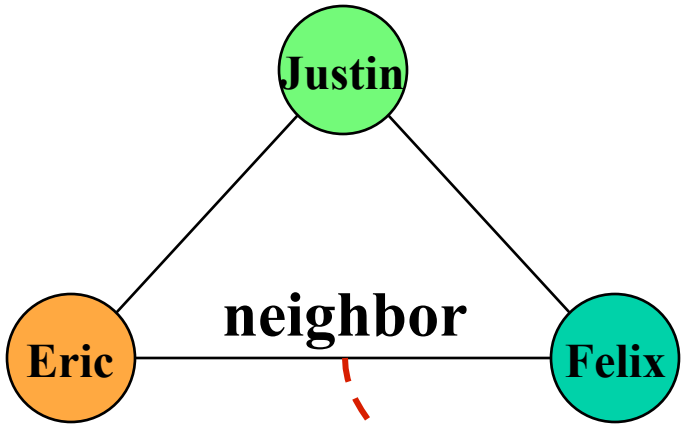

FB friends: 790+
FV neighbors: 30+

Is this necessary?

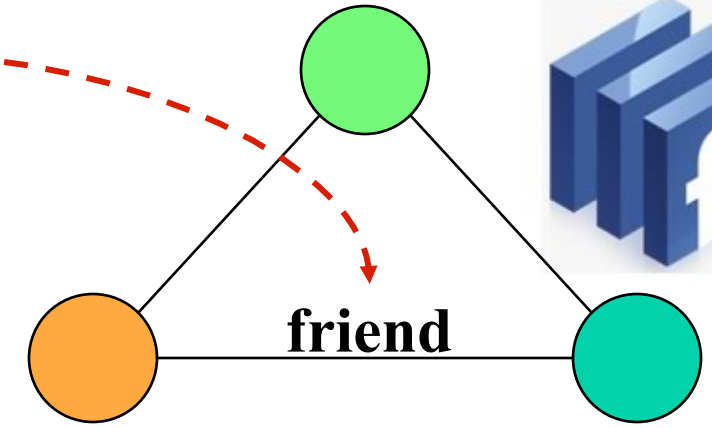

Social Network Transformation

- What is the best/effective VPSN for this application X?
  - Virtual Private Social Network

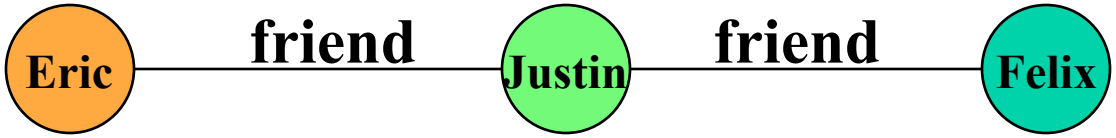

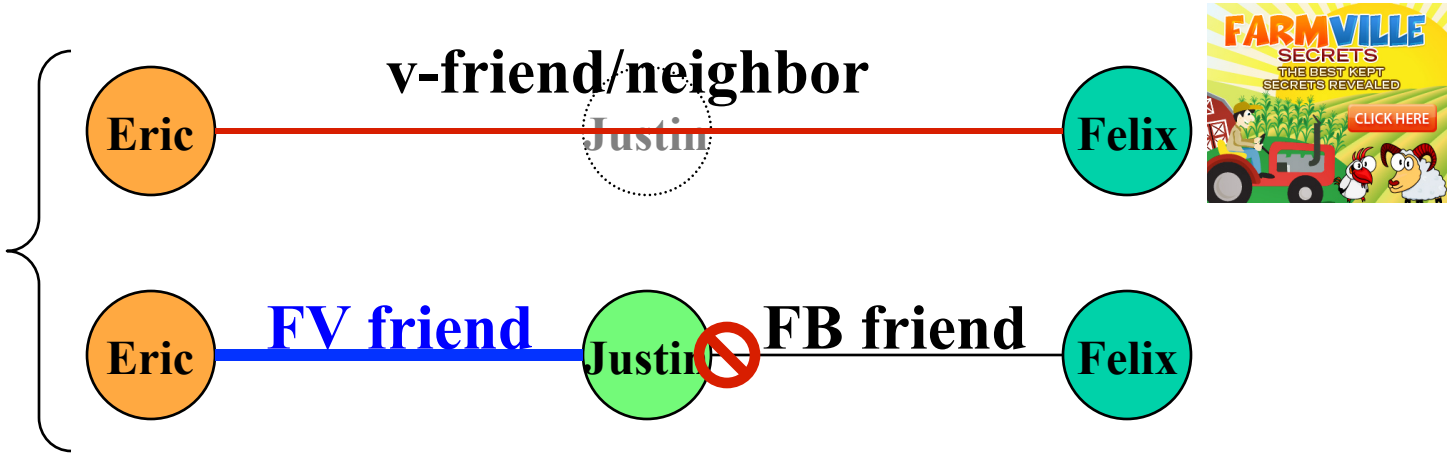

Social Network Transformation

- Each application might need different OSN topologies.

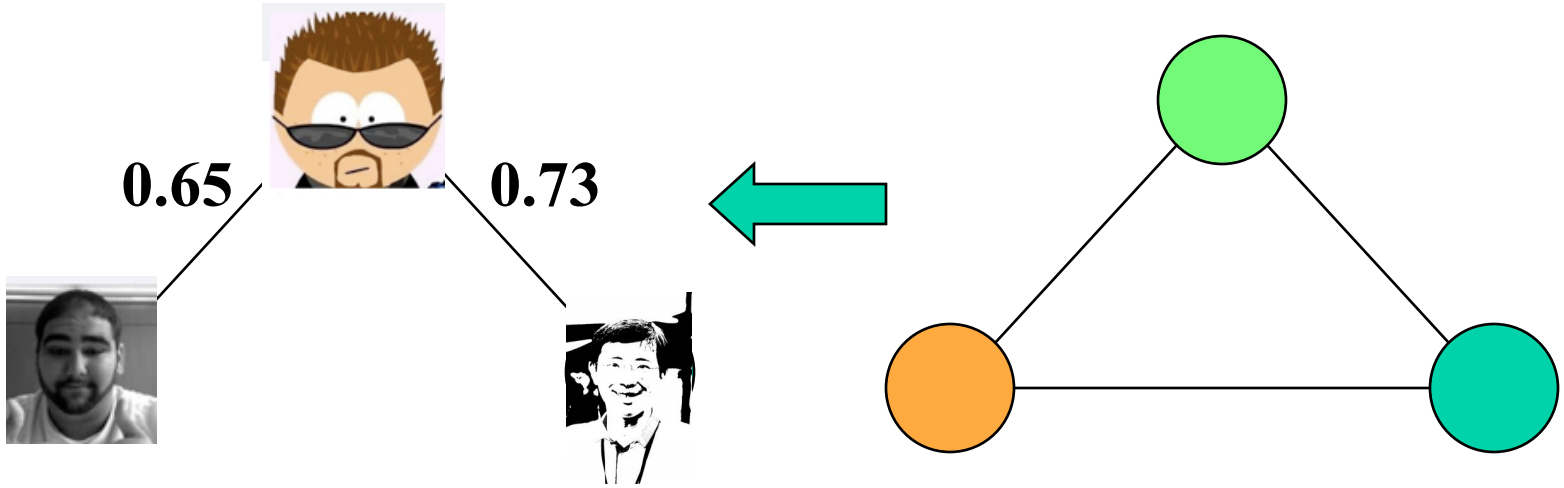

Social-Centric OS

- Managing Social Resources

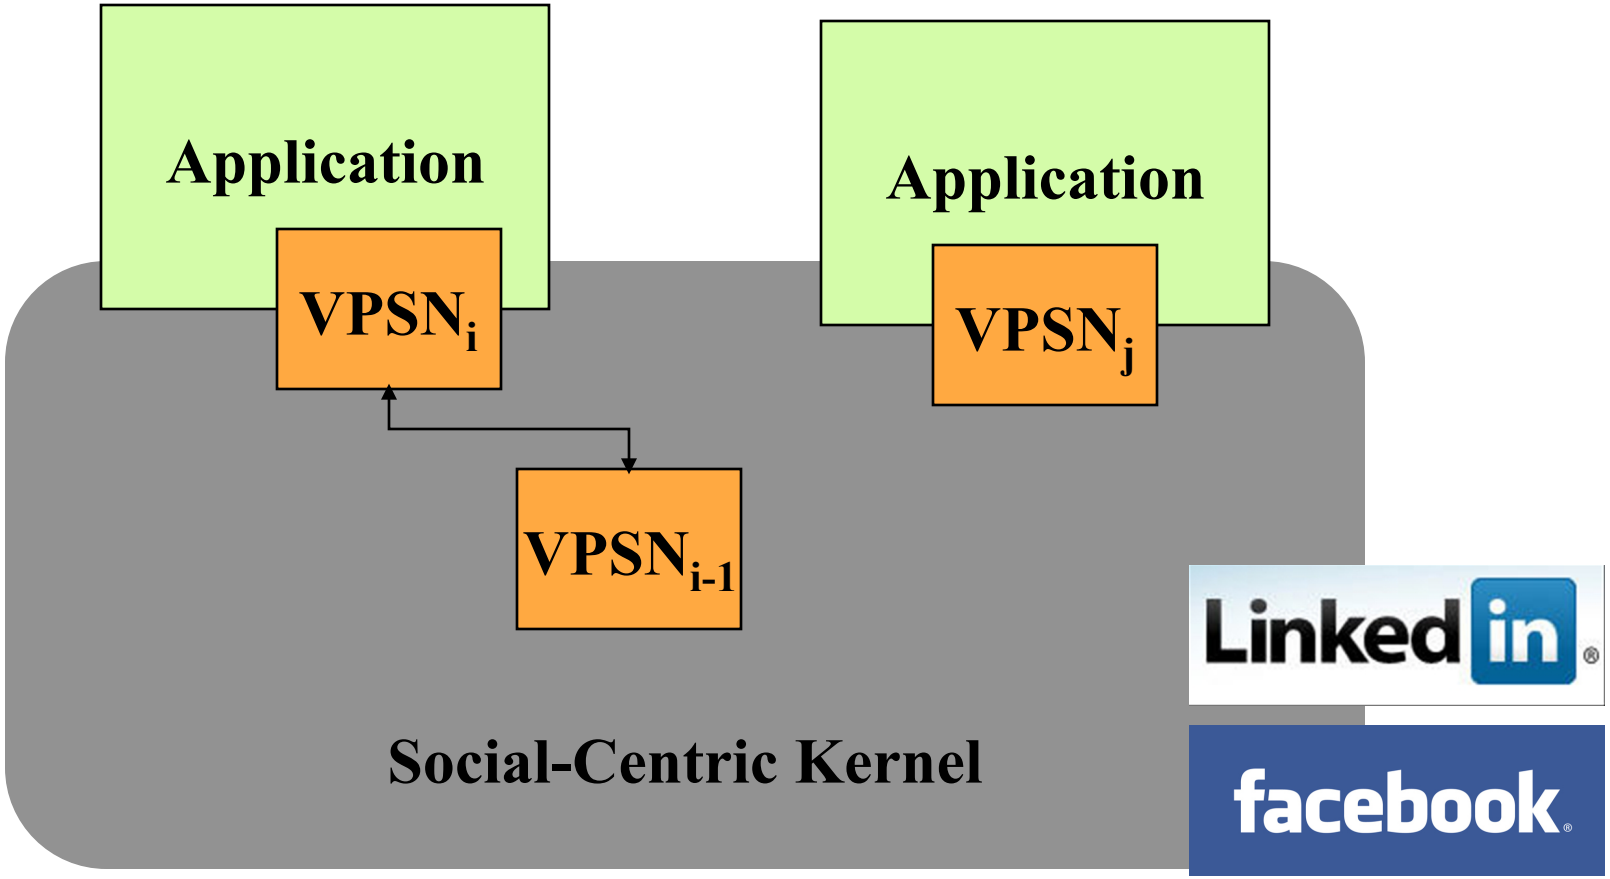

Social Network Transformation

UC DAVIS
UNIVERSITY OF CALIFORNIA

- Let "users" (or community of users) define and control their social needs.
- And, the impact of VPSN to the backbone OSN or other VPSN needs to be carefully managed. (open problem)
  - Currently, we only provide "CSI isolation"

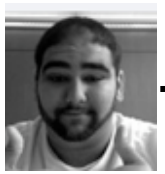

Application

Existing Applications

Social-Enabled Applications and Games

Name-ID resolution

Community Oriented Keywords

Policy/Reputation-based Route discovery

Wrapper

Social Context

SCOSK

OSN

Felix

Eric

tagging

Social network transformation

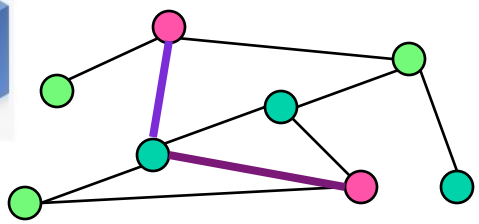

SCOSK over OSN

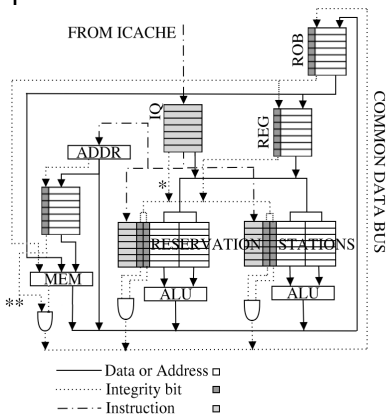

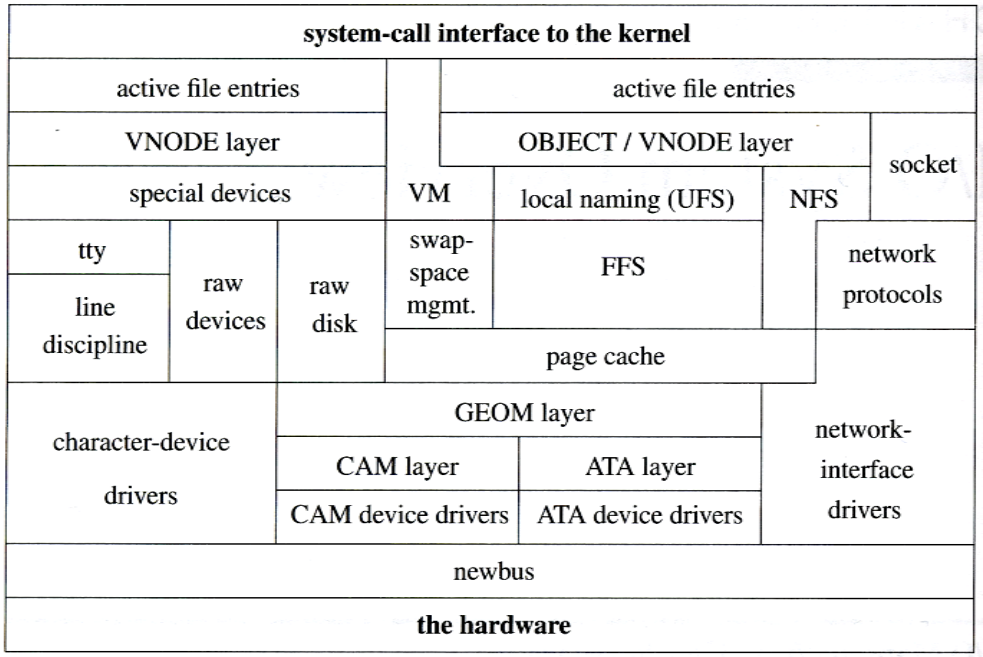

Social Control System Calls

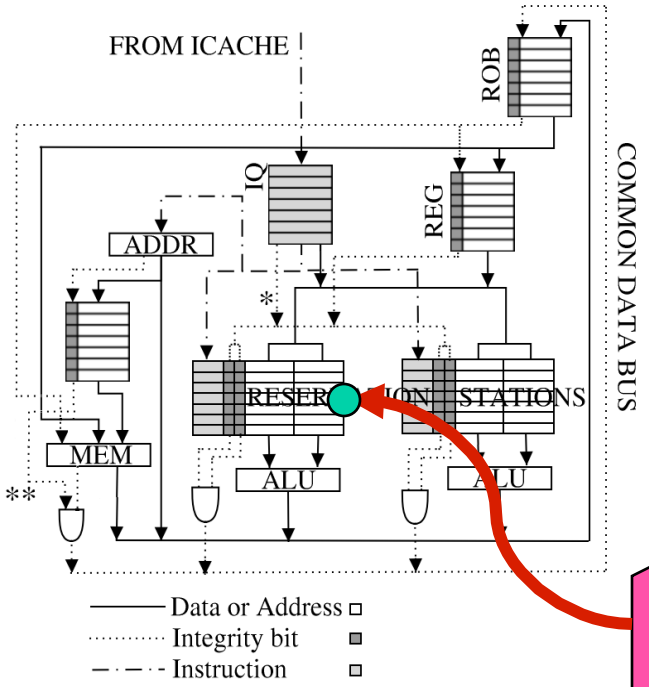


```
int CSI(…);
int propagate(vnode *);
CSI *get(vnode *);
int wait(vnode *, cond_t *);
int signal(vnode *, cond_t *);
```

X-DSL 0x15EF2AC4

NE2000

How many? within how much time?

UCDAVIS UNIVERSITY OF CALIFORNIA

- Social capitals/resources

"anomalous" social transactions

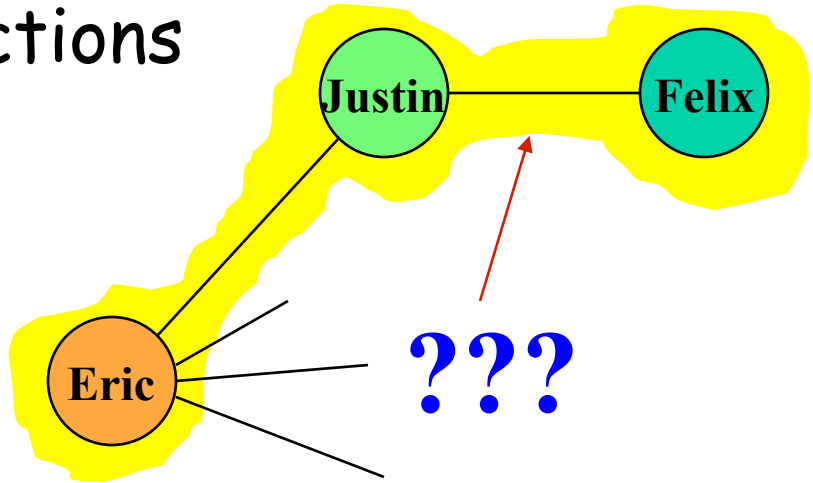

The response from the dark side….

- **obtain the expected social profile**
- **leverage a large number of compromised social bots**
- **each produce/consume a very small amount of SR**

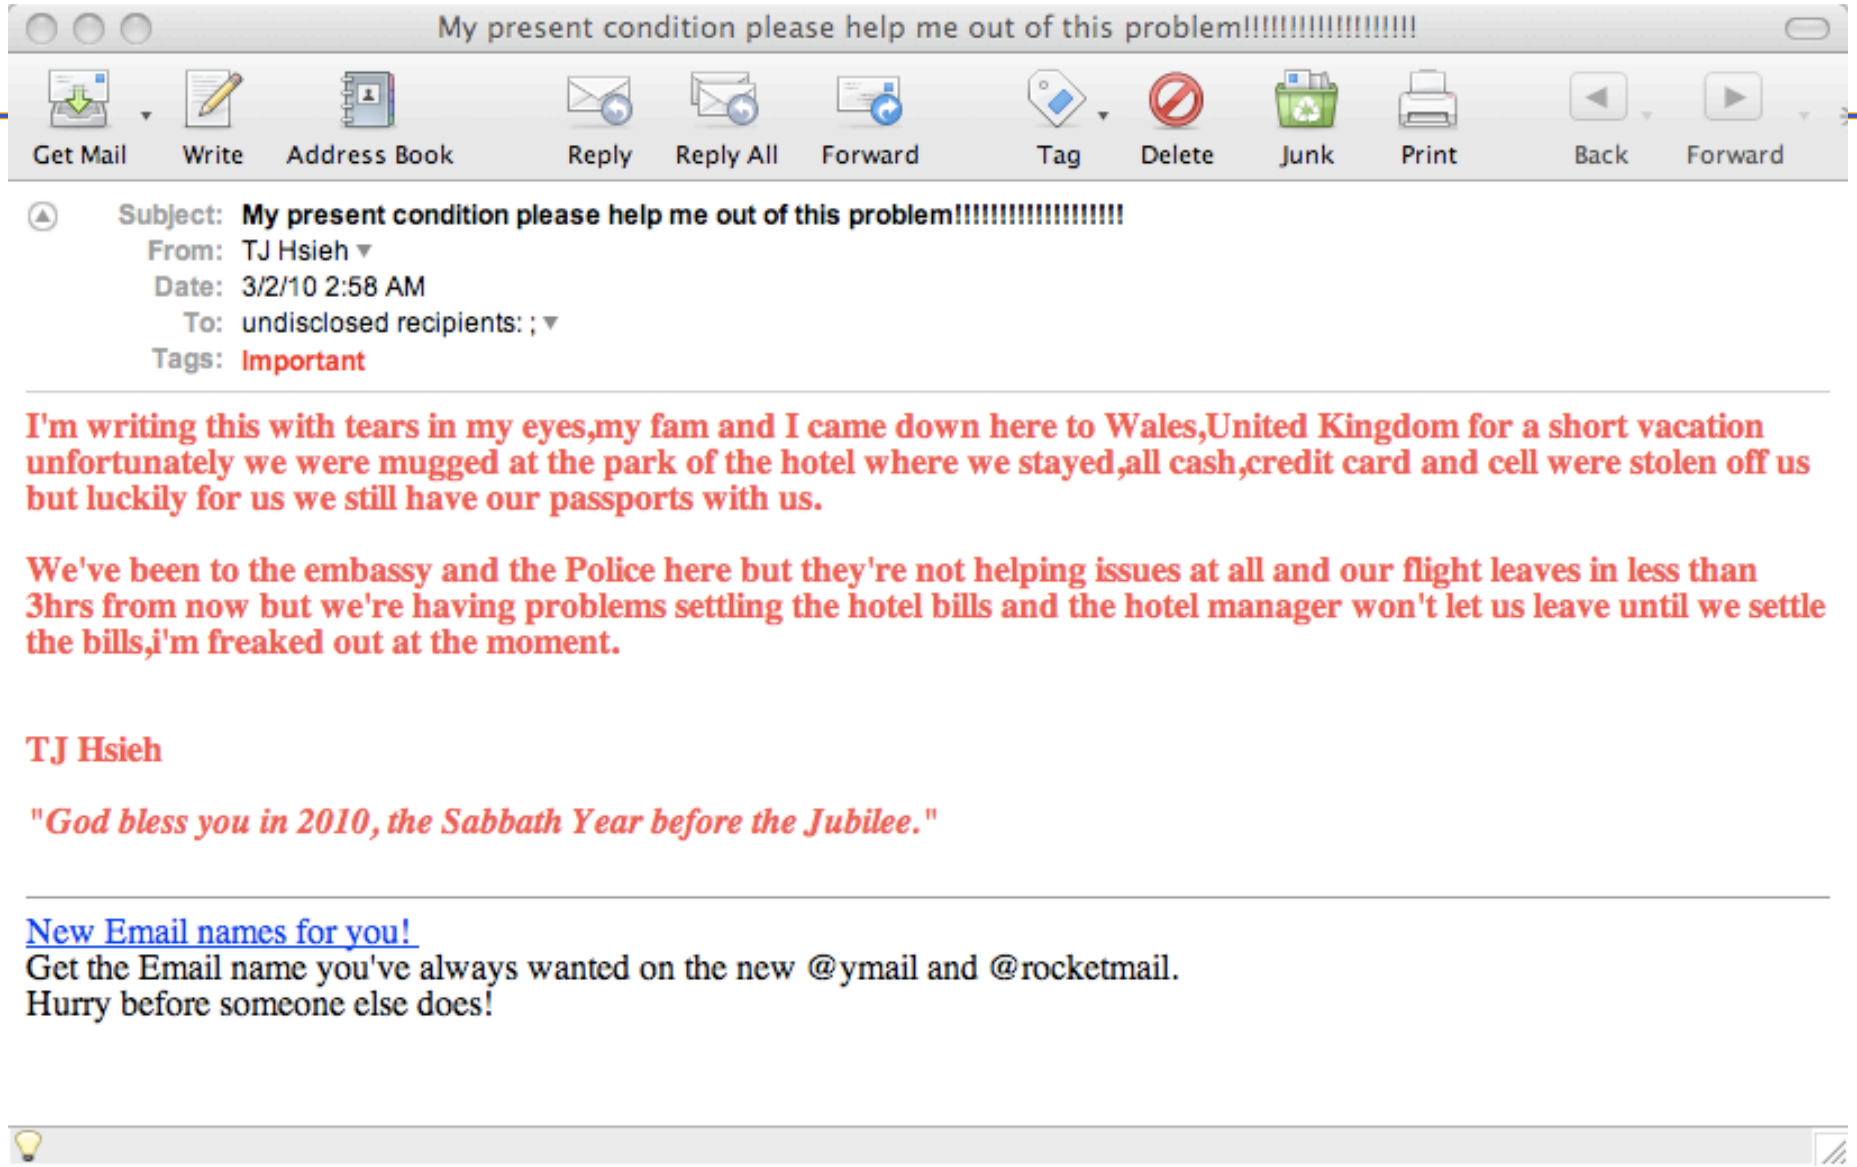
My present condition please help me out of this problem!!!!!!!!!!!!!!!!!!!!
Get Mail
Write
Address Book
Reply
Reply All
Forward
Tag
Delete
Junk
Print
Back
Forward
Subject: My present condition please help me out of this problem!!!!!!!!!!!!!!!!!!!!
From: TJ Hsieh
Date: 3/2/10 2:58 AM
To: undisclosed recipients: ;
Tags: Important
I'm writing this with tears in my eyes,my fam and I came down here to Wales,United Kingdom for a short vacation unfortunately we were mugged at the park of the hotel where we stayed,all cash,credit card and cell were stolen off us but luckily for us we still have our passports with us.
We've been to the embassy and the Police here but they're not helping issues at all and our flight leaves in less than 3hrs from now but we're having problems settling the hotel bills and the hotel manager won't let us leave until we settle the bills,i'm freaked out at the moment.
TJ Hsieh
"God bless you in 2010, the Sabbath Year before the Jubilee."
New Email names for you!
Get the Email name you've always wanted on the new @ymail and @rocketmail.
Hurry before someone else does!

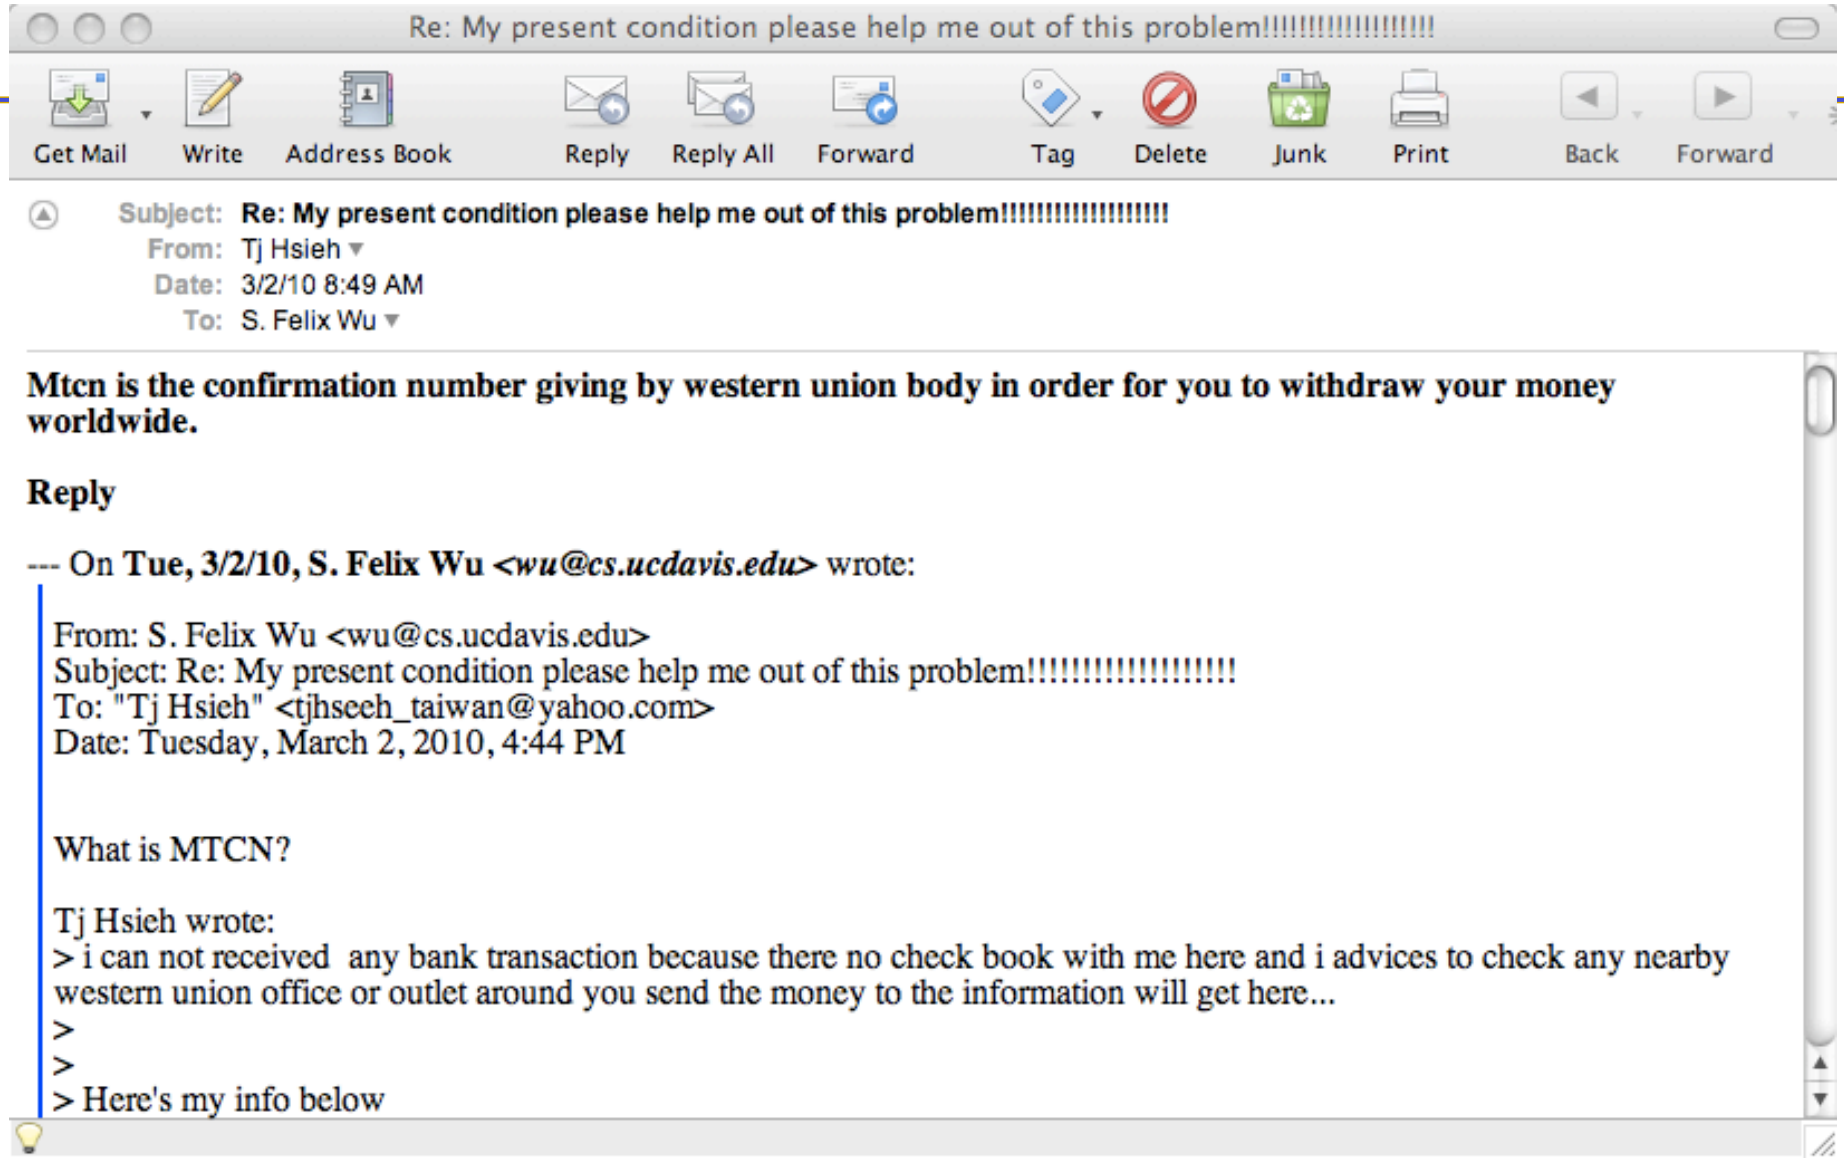
Re: My present condition please help me out of this problem!!!!!!!!!!!!!!!!!!!
Get Mail
Write
Address Book
Reply
Reply All
Forward
Tag
Delete
Junk
Print
Back
Forward
Subject: Re: My present condition please help me out of this problem!!!!!!!!!!!!!!!!!!!
From: Tj Hsieh
Date: 3/2/10 8:49 AM
To: S. Felix Wu
Mtcn is the confirmation number giving by western union body in order for you to withdraw your money worldwide.
Reply
--- On Tue, 3/2/10, S. Felix Wu <wu@cs.ucdavis.edu> wrote:
From: S. Felix Wu <wu@cs.ucdavis.edu>
Subject: Re: My present condition please help me out of this problem!!!!!!!!!!!!!!!!!!!
To: "Tj Hsieh" <tjhseeh_taiwan@yahoo.com>
Date: Tuesday, March 2, 2010, 4:44 PM
What is MTCN?
Tj Hsieh wrote:
> i can not received any bank transaction because there no check book with me here and i advices to check any nearby western union office or outlet around you send the money to the information will get here...
>
>
> Here's my info below

How to build the trust?

- If only I have some social relationship information available…

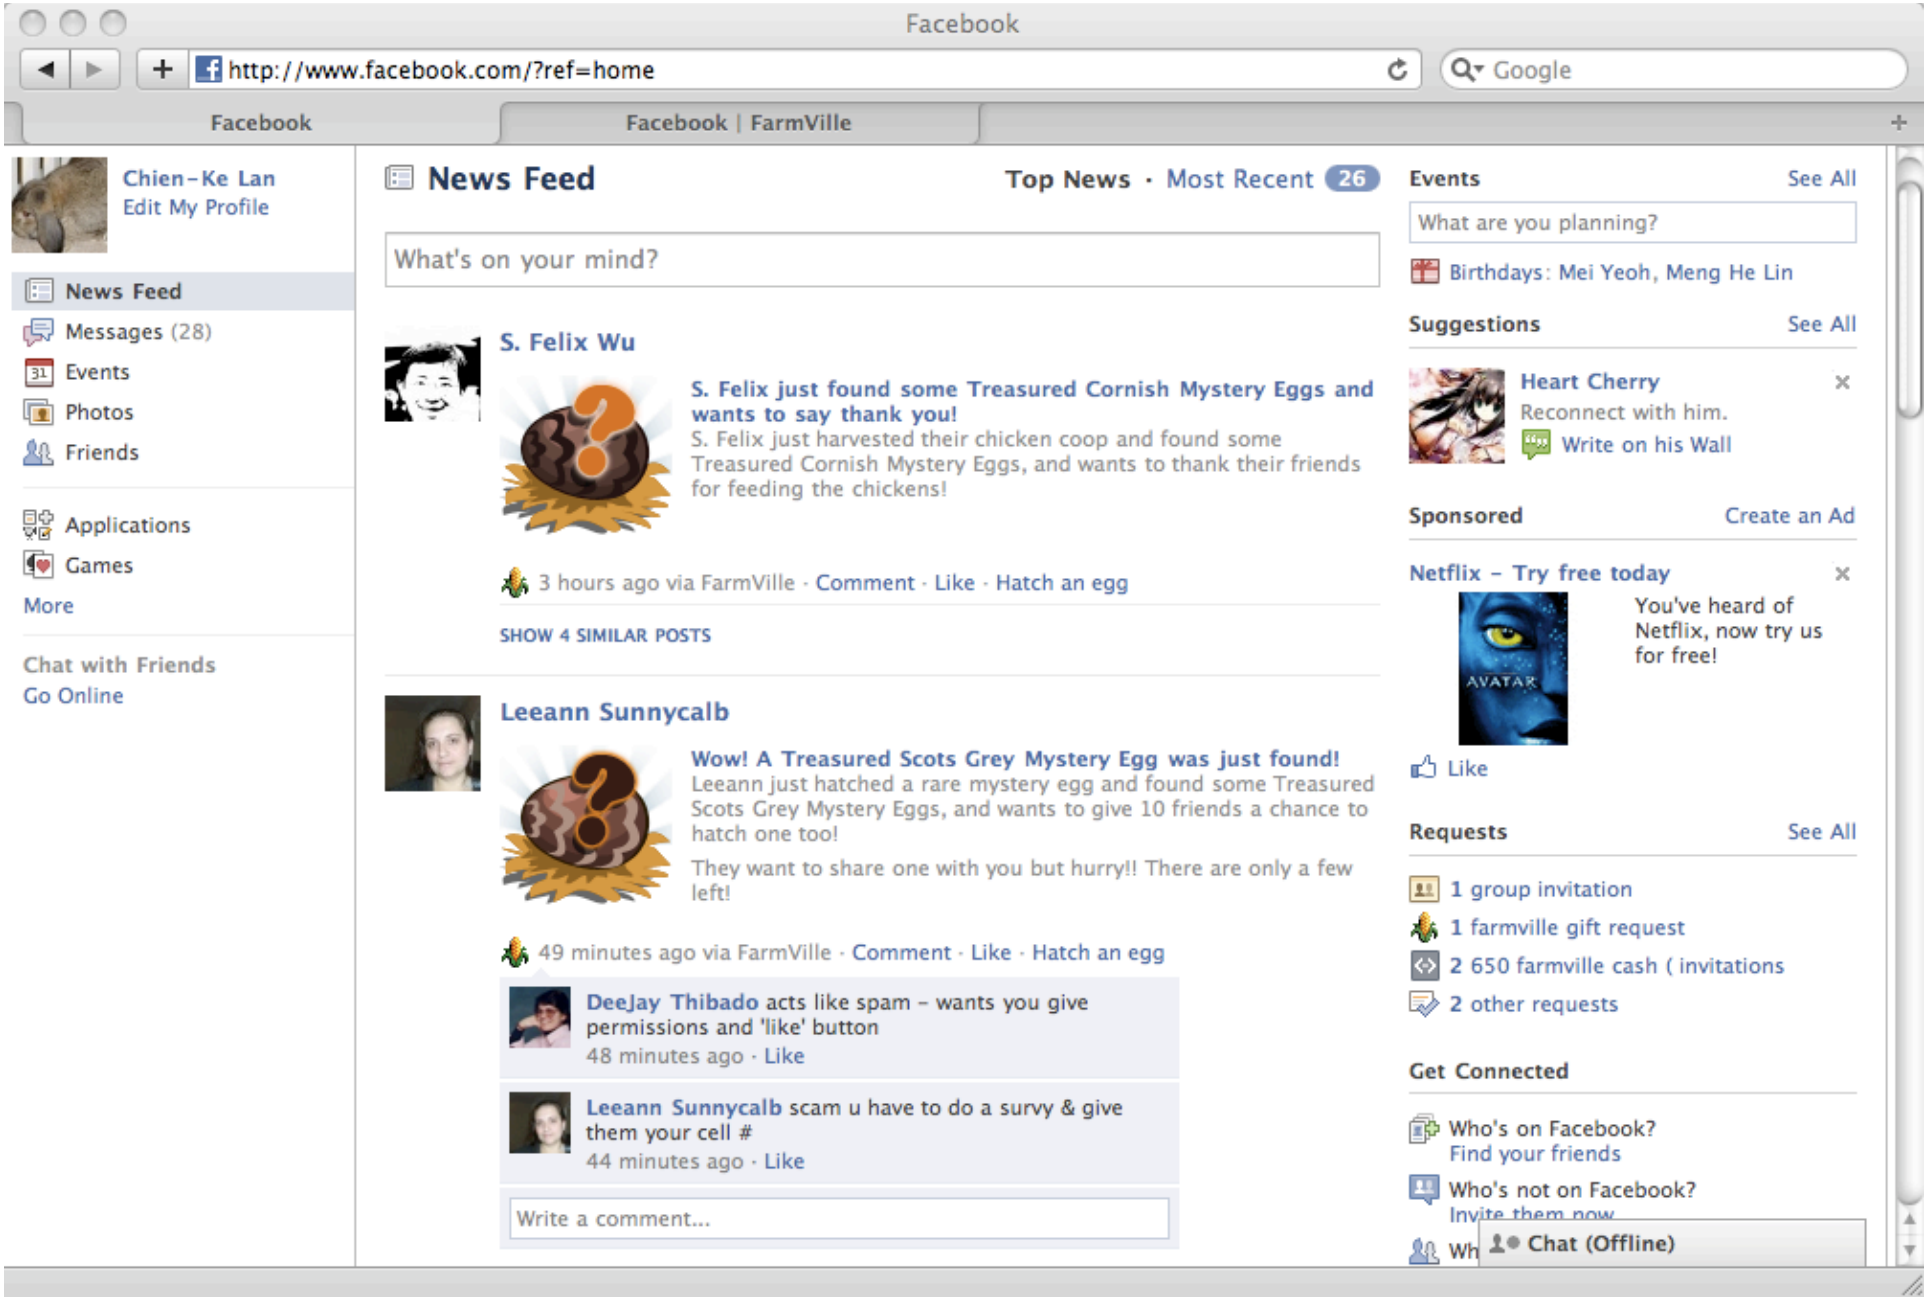
Facebook
http://www.facebook.com/?ref=home
Google
Facebook
Facebook | FarmVille
Chien-Ke Lan
Edit My Profile
News Feed
Messages (28)
Events
Photos
Friends
Applications
Games
More
Chat with Friends
Go Online
News Feed
Top News · Most Recent 26
What's on your mind?
S. Felix Wu
S. Felix just found some Treasured Cornish Mystery Eggs and wants to say thank you!
S. Felix just harvested their chicken coop and found some Treasured Cornish Mystery Eggs, and wants to thank their friends for feeding the chickens!
3 hours ago via FarmVille · Comment · Like · Hatch an egg
SHOW 4 SIMILAR POSTS
Leeann Sunnycalb
Wow! A Treasured Scots Grey Mystery Egg was just found!
Leeann just hatched a rare mystery egg and found some Treasured Scots Grey Mystery Eggs, and wants to give 10 friends a chance to hatch one too!
They want to share one with you but hurry!! There are only a few left!
49 minutes ago via FarmVille · Comment · Like · Hatch an egg
DeeJay Thibado acts like spam – wants you give permissions and 'like' button
48 minutes ago · Like
Leeann Sunnycalb scam u have to do a survy & give them your cell #
44 minutes ago · Like
Write a comment...
Events
See All
What are you planning?
Birthdays: Mei Yeoh, Meng He Lin
Suggestions
See All
Heart Cherry
Reconnect with him.
Write on his Wall
Sponsored
Create an Ad
Netflix – Try free today
You've heard of Netflix, now try us for free!
Like
Requests
See All
1 group invitation
1 farmville gift request
2 650 farmville cash (invitations
2 other requests
Get Connected
Who's on Facebook?
Find your friends
Who's not on Facebook?
Chat (Offline)


Facebook | Chien-Ke Lan
http://www.facebook.com/profile.php?id=100000880093751
Google
Facebook | Chien-Ke Lan
Facebook | FarmVille
Chien-Ke Lan
Wall
Info
Photos
綁匪之王
義氣仔女Online
Notes
Edit My Profile
FBR (Facebook Ranger)
Information
Relationship Status:
It's Complicated
Birthday:
June 30, 1971
Friends
541 friends
See All
吳忠彥
吳子青
Lee Han Kim
孫君瑞
Rainbow King
Wai Hung Yeung
What's on your mind?
Attach:
Share
Options
Chien-Ke Lan
Wow! A Treasured Scots Grey Mystery Egg was just found!
Chien-Ke just hatched a rare mystery egg and found some Treasured Scots Grey Mystery Eggs, and wants to give 10 friends a chance to hatch one too!
They want to share one with you but hurry!! There are only a few left!
6 minutes ago via FarmVille · Comment · Like · Hatch an egg
Chien-Ke Lan
Chien-Ke was working in the stables!
Chien-Ke's horses were so helpful on their farm, Chien-Ke was able to get Arborists to help them! Ch...
See More
about an hour ago via FarmVille · Comment · Like · Get Help on Your Farm
Chien-Ke Lan
Remove
Chien-Ke just found some Treasured Golden Mystery Eggs and wants to say thank you!
Chien-Ke just harvested their chicken coop and found some Treasured Golden Mystery Eggs, and wants to thank their friends for feeding the chickens!
about an hour ago via FarmVille · Comment · Like · Hatch an egg
Create an Ad
Connect in more places
at&t
FastAccess® DSL Lite with Wi-Fi included keeps you connected in more places. Act now and get 3 months Free. $19.95/mo. Order now.
Like
Sacramento Almost Free
75% off smoothies? A one hour massage for $30? A sunset cruise for $20? See today's Sacramento deal!
Like
iPad up to 75% off
Chat (Offline)
Display a menu for "http://www.facebook.com/apps/application.php?id=130529466968382"

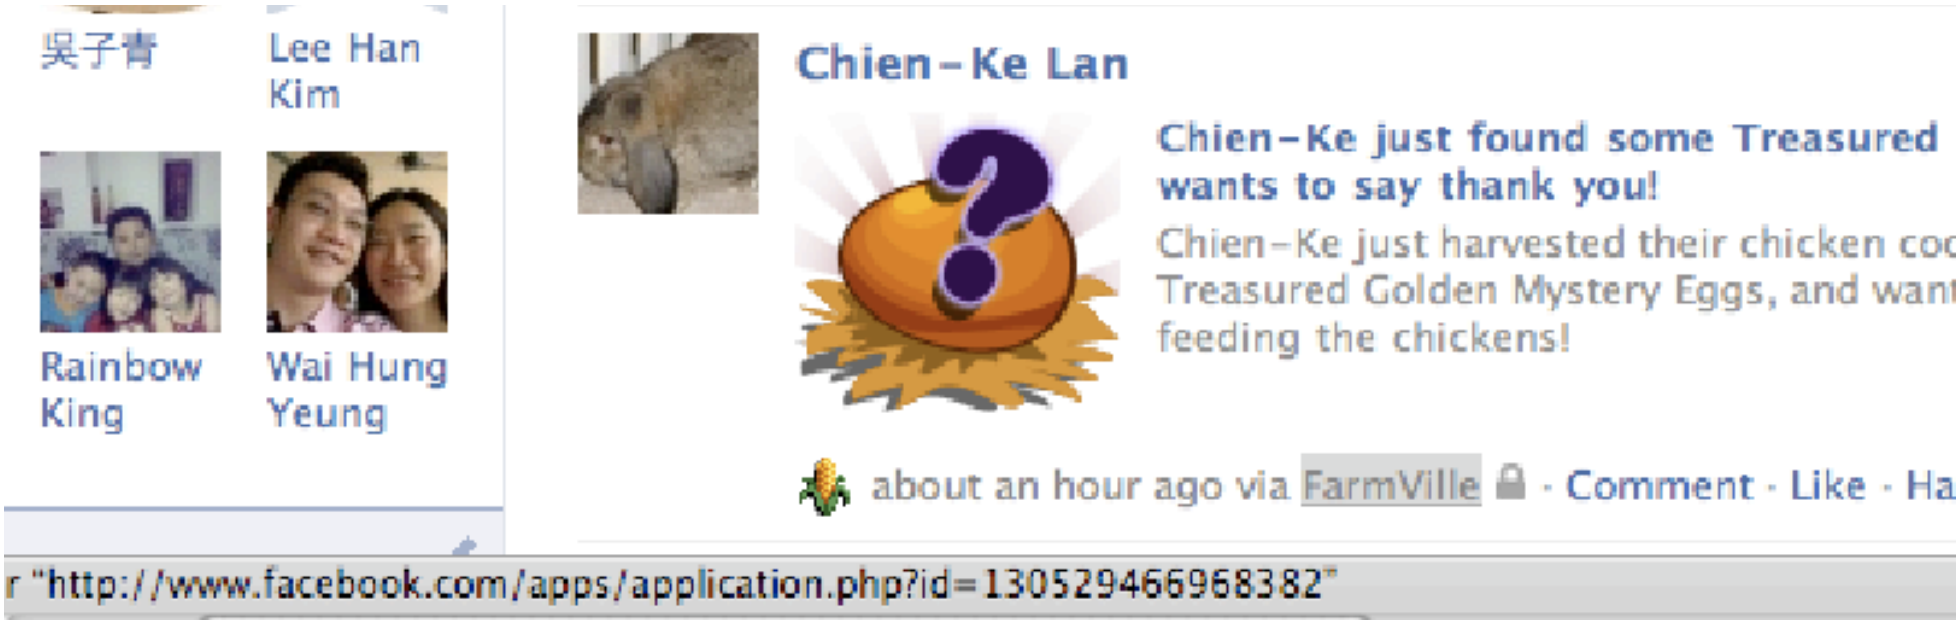
吳子青
Lee Han Kim
Rainbow King
Wai Hung Yeung
Chien-Ke Lan
Chien-Ke just found some Treasured
wants to say thank you!
Chien-Ke just harvested their chicken coo
Treasured Golden Mystery Eggs, and want
feeding the chickens!
about an hour ago via FarmVille · Comment · Like · Ha
r "http://www.facebook.com/apps/application.php?id=130529466968382"


Facebook | FarmVille
http://www.facebook.com/apps/application.php?id=130529466968382&v=info
RSS
Google
Facebook | Chien-Ke Lan
Facebook | FarmVille
facebook
Search
FarmVille
Like
Wall
Info
Reviews
Discussions
Go to Application
Add to my Page
Suggest to Friends
Block Application
About FarmVille
Friends using this Application
No friends are using this application.
About the Developer
Bill Barrelli
Information
Category
All
This application was not developed by Facebook.
Chat (Offl

Gamers Unite!

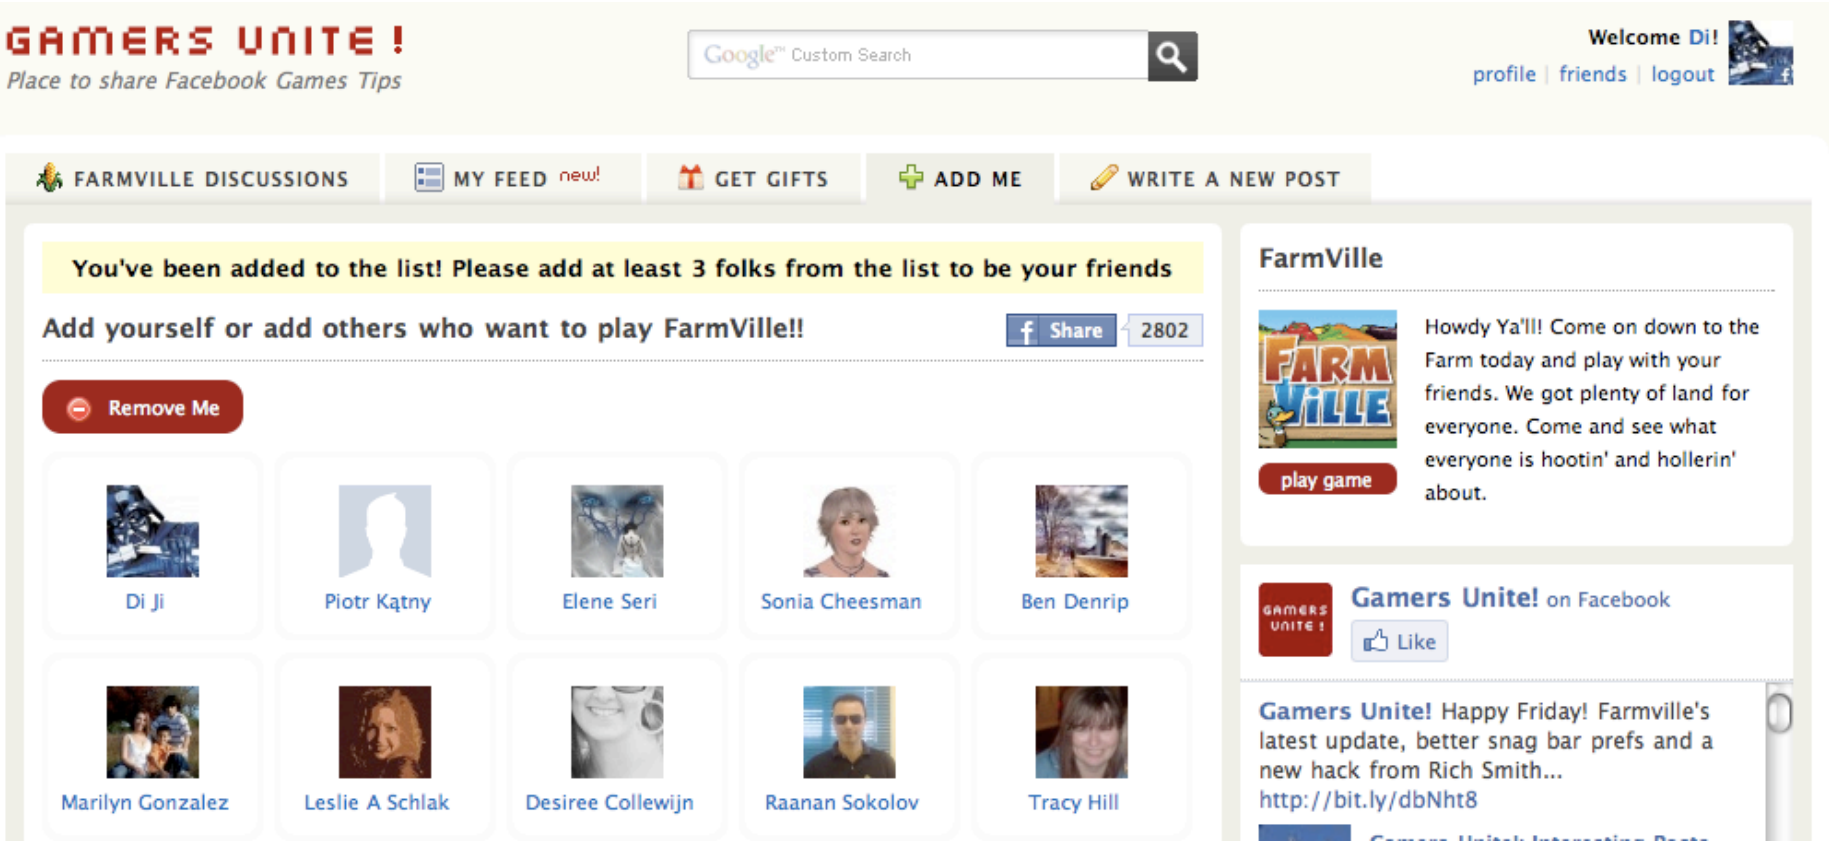

How to build the trust?

- If only I have some social relationship information available…

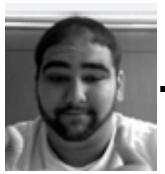

Application

Existing Applications

Social-Enabled Applications and Games

Name-ID resolution

Community Oriented Keywords

Policy/Reputation-based Route discovery

Wrapper

Social Context

SCOSK

OSN

Felix

Eric

tagging

Social network transformation

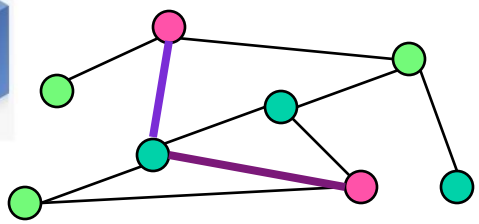

SCOSK over OSN

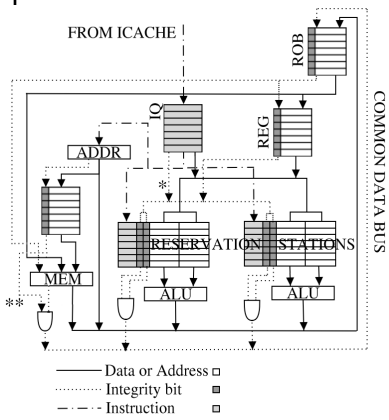

UCDAVIS
UNIVERSITY OF CALIFORNIA

DEMO time…

Understanding Dynamics of Behavior in Social Circles

Jennifer Golbeck and Ugur Kuter
University of Maryland, College Park

CSI: College Park

- Research at the intersections of social trust networks, probabilistic reasoning, game theory and planning

Game theory

Probabilistic Models

Social Trust

Web Service Composition

Social Networks

- Web Page (under construction): www.csicp.org

In this talk,

- How to make decisions and act, when
  - you're working on a physical transportation network,
  - you're interacting with a social group,
  - you do not know individual tendencies (e.g., altrusim, fairness, individualistic) in the group,
  - you do not know the group's collective tendencies (e.g., cooperativeness).

Outline

- Understanding cooperativeness in groups
  - Cooperativeness as a measure of intentions of the group
  - Other factors that go into intentions/objectives understanding?
- Ongoing: Learning how to adapt to an adversarial group
- Ongoing: Dynamics of social trust relationships among a group of agents
  - Among humans as well as mixed human/software agent groups
  - Evaluation framework: strategy games
- Conclusions and Future Work

Understanding Cooperativeness

- In most of the time, we're living in a Darwinian society where the survival of the fittest is the prominent rule

- Research Questions:
  - How and why cooperative behavior emerges given that selfish individuals can gain by taking advantage of other?
  - How to model an agent behavior in a social environment?
  - How to define the ***cooperativeness*** of a group?

Evolution of Cooperation

- Existing research typically utilize games for modeling social dilemmas about cooperating and defecting
  - E.g., Prisoner's Dilemma (PD)

Prisoner's Dilemma		Player 2	
		Cooperate (C)	Defect (D)
Player 1	Cooperate (C)	(3, 3)	(0, 5)
	Defect (D)	(5, 0)	(1, 1)

  - Played only once:
    - Defect (D) always provides a higher individual payoff
    - But mutual cooperation (C) is socially optimal
  - Played **repeatedly** for an unknown number of times:
    - i.e., Iterated Prisoner's Dilemma (IPD)
    - Cooperative behavior (e.g., Tit-for-Tat) might emerge to increase accumulated payoff

Social Value Orientation (SVO) Theory

- SVO is a social choice theory [Messick and McClintock in 1968]

- Based on empirical evidence, SVO states that:
  - Individuals **vary in their interpersonal social tendencies**
    - Individual's behavior may often be **motivated** not only by self-interest but **also by the consequences for the others**
  - These reflect **stable** and **consistent** orientations

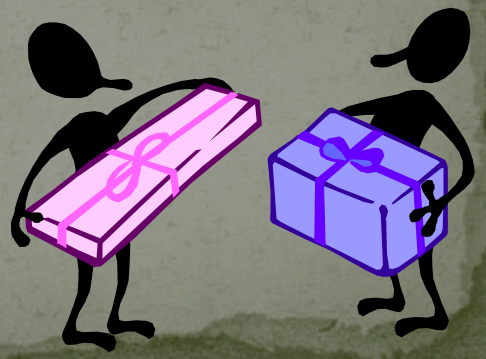

Our Contributions (GECCO-10, IAT-10)

- Proposed a new model
  - Based on **Social Value Orientation** (SVO) theory in social and behavioral science
  - Explicitly **captures the notion of pro-social vs. pro-self** orientations exhibited in human behavior
- Analyzed steady state behaviors
- Provided a new definition for cooperativeness using social orientation of players
  - Independent of the average payoff and game matrix
  - Accurately describe some situations that using average payoffs as a measure of cooperativeness cannot

Our SVO-inspired Model

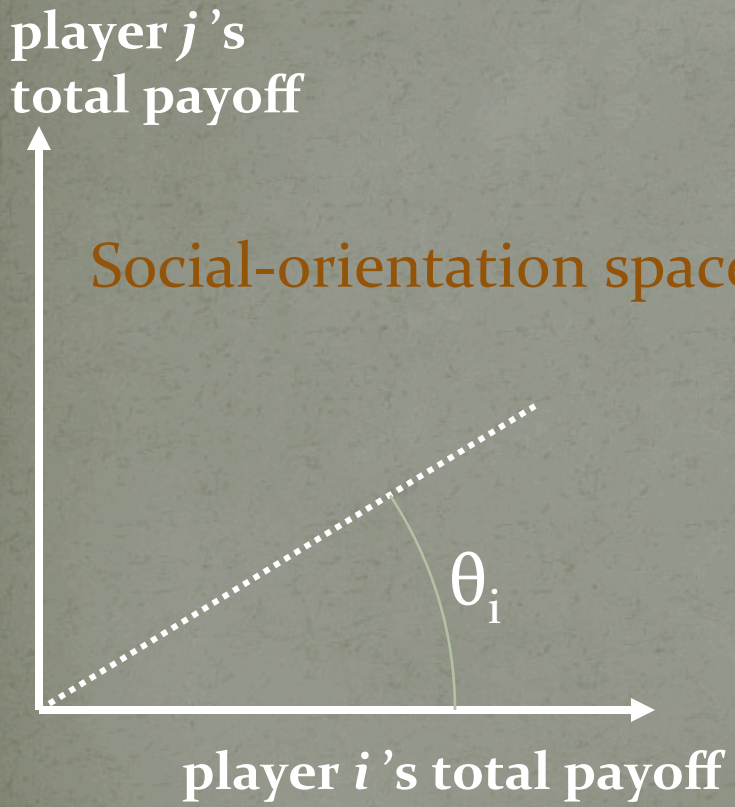

- Social orientation (θ_i) of Player i
  - A model of the player's tendency to adopt a pro-social or pro-self behavior
- Examples:
  - $\theta_i = 0$ => individualistic
  - $\theta_i = \pi/4$ => fair
  - $\theta_i = \pi/2$ => altruistic

How a player behaves based on its θ?

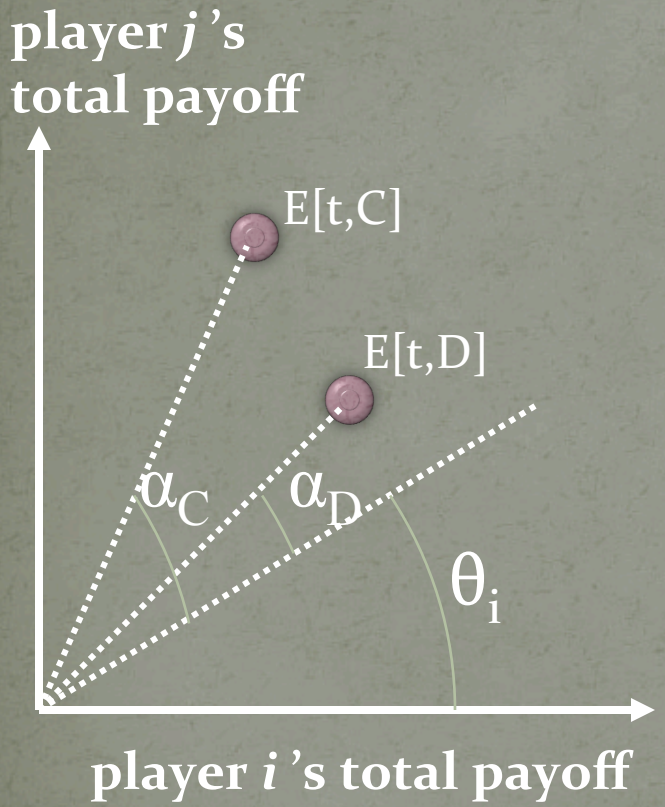

- Each player aims to bring its own expected payoff closer to its social-orientation
  - E(t,D): accumulated payoff that Player i receives until iteration t
  - Choose an action, a, such that the deviation angle α_a is minimized

$$a = \text{argmax}_{a \in A} \cos \alpha_a$$

Summary of results

2x2 symmetric game		Player 2	
		Cooperate (C)	Defect (D)
Player 1	Cooperate (C)	(R, R)	(S, T)
	Defect (D)	(T, S)	(P, P)

(WLOG, assume $\theta_i \leq \theta_j$)	$\theta_i + \theta_j < 90$	$\theta_i + \theta_j = 90$	$\theta_i + \theta_j > 90$
Both players are pro-self $\Theta_i < 45, \theta_j < 45$			Both players always defect at each game in steady state
Both players are pro-social $\Theta_i > 45, \theta_j > 45$	Both players always cooperate at each game in steady state		
One pro-self, one prosocial, i.e., $\theta_i < 45, \theta_j > 45$	Pro-self player i gets $\Theta_i(T+S)$, and pro-social player j gets $(1-\Theta_i)(T+S)$ at each game in steady state	Player 1 gets $\bar{p}_i = \frac{SP-PT}{(P-T)-(P-S)\frac{1-r_i}{r_i}}$ Player 2 gets $\bar{p}_j = \bar{p}_i \frac{1-r_i}{r_i}$	Player 1 gets $\bar{p}_i = \bar{p}_j \frac{1-r_j}{r_j}$ Player 2 gets $\bar{p}_j = \frac{TR-RS}{(R-S)-(R-T)\frac{1-r_j}{r_j}}$

How to Define the Cooperativeness?

- Traditional definition: Average payoff obtained by players
  - Higher the average payoff, more cooperative the society
  - May not be accurate, e.g., defecting may lead to high payoff sometimes
    - We will illustrate it in the following experiments

- Our definition: Social orientation of players
  - $\theta_i \geq \pi/4$ => Player i is cooperative
  - Cooperativeness of a society = avg θ or % of cooperative agents
  - Independent of the average payoff and game matrix

Experiments

- Investigate emergence of cooperative populations based on social orientations of individuals

- Use Evolutionary Simulations
  - Players play IPD with each other
  - Replicator dynamics
    - each player has a number of offspring that is proportional to its expected total payoff
  - Initial population
    - Randomly generated 10 θ values from the interval $[0, \pi/2]$
    - Size of a group with a particular θ value constitutes 10% of the entire population
  - Mutation
    - On average, once in every 100 generations, introduce a small amount of new randomly-chosen mutant players

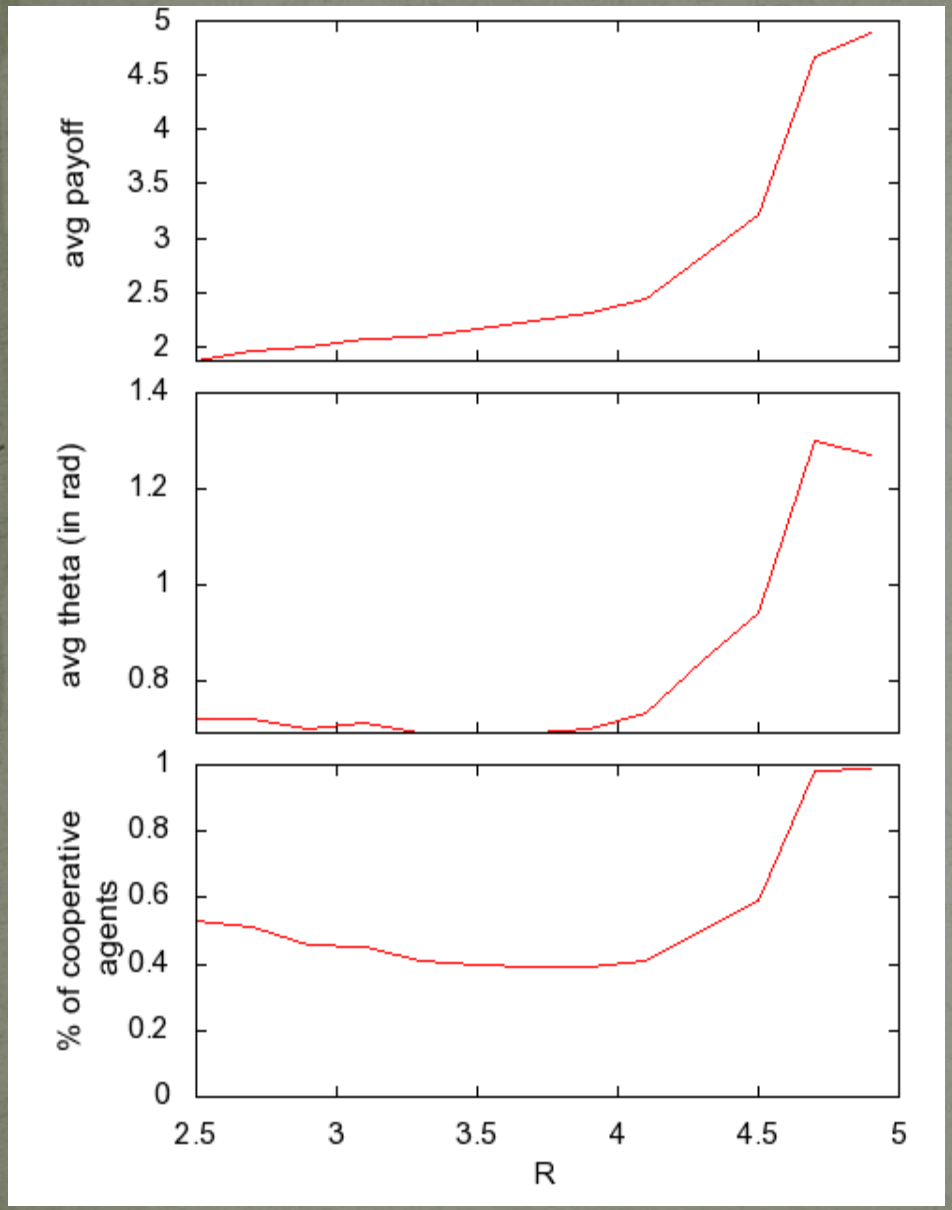

Varying R

2x2 symmetric game		Player 2	
		Cooperate (C)	Defect (D)
Player 1	Cooperate (C)	(R, R)	(S, T)
	Defect (D)	(T, S)	(P, P)

- R in [2.5, 5]
  - keeping the S, T and P constant with their original values
  - keeping the preference relations in the PD matrix, i.e., S < P < R < T and 2R > S + T
- Result: pro-social tendency increases with increasing reward (R) as expected

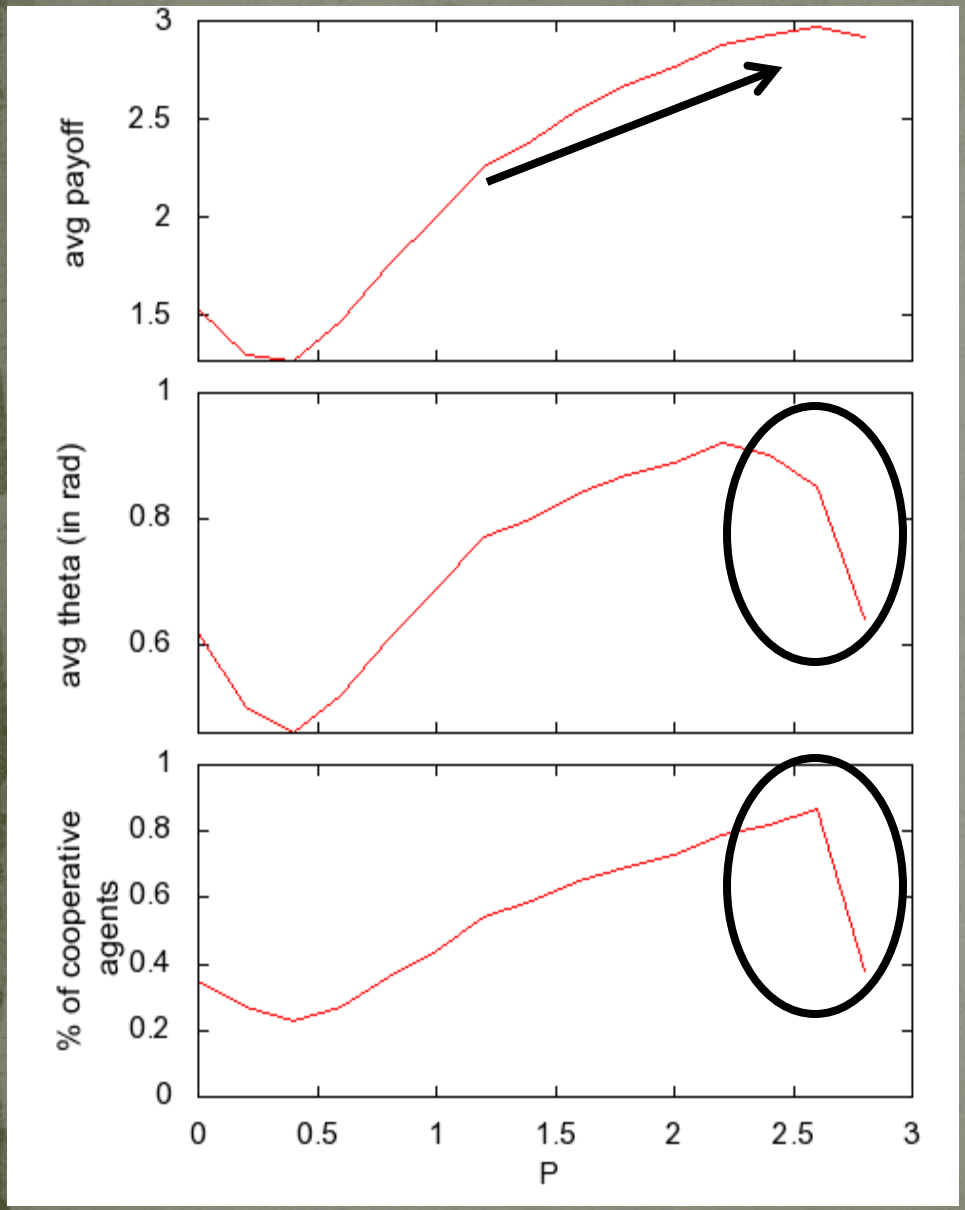

Varying P

2x2 symmetric game		Player 2	
		Cooperate (C)	Defect (D)
Player 1	Cooperate (C)	(R, R)	(S, T)
	Defect (D)	(T, S)	(P, P)

- Average payoff increases when P increases
- Average θ drops sharply when P is large
  - Contradicts with the predictions of standard definition of cooperativeness!
    - No one-to-one correlation between the average payoff and the group's cooperativeness

Why are the evolved groups different?

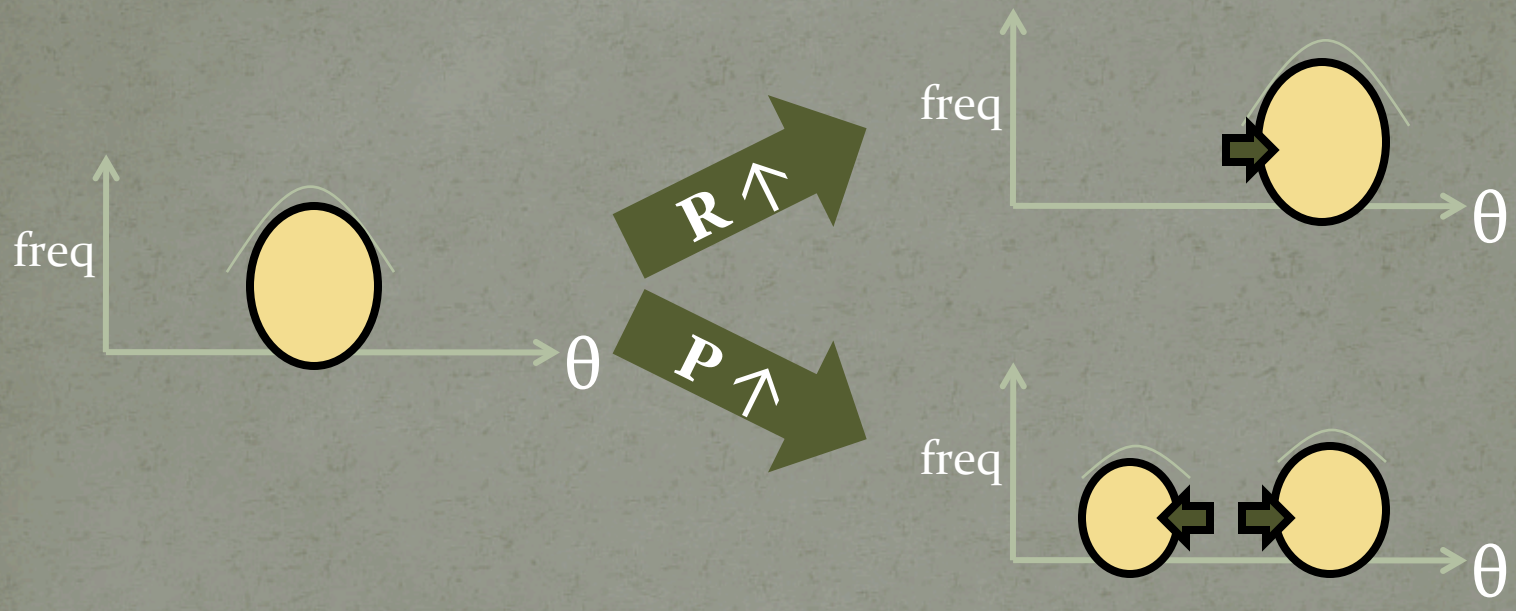

Importance of this work

- Showed the advantage of using social orientation of players and payoffs as a metric for the society's cooperativeness level

- However, this result suggests that understanding, game-theoretically, the dynamics of a group's behavior has some open, interesting issues
  - While previous works using average payoff as a metric of the degree of cooperation may not be accurate sometimes
  - Next slides on two ongoing works on social trust in groups

Trust and Credibility in Social Groups

- We will define and implement a series of increasingly complex games that involve
  - information-gathering, information-hiding, negotiations over physical resources
- In developing these games, we will require the following features:
  - Players do not know who is working with them or against them, so trust is initially unknown.
  - Players must interact and exchange resources (e.g. money, information, etc.) so there are bases for establishing a history that will influence trust.
  - Payoffs will depend on identifying who is trustworthy, so players have an incentive to investigate and establish trust in others.
- Objective: analyze the game data to understand how the players develop trust in each other and how their trust changes over the course of the games

- There are three good players and one bad player.
- All players are heading to the center.
  - The bad player wins if he gets to the center ahead of at least one good player.
  - Otherwise the good team wins.
- Players draw cards at the beginning of the game that will state if they are good or bad.
- Resources and requirements
  - Resources can be exchanged between players.
  - They can be traded, given away, or exchanged with promise of something in return later.
  - There are no rules to prevent players from lying or going back on promises - anything is fair!

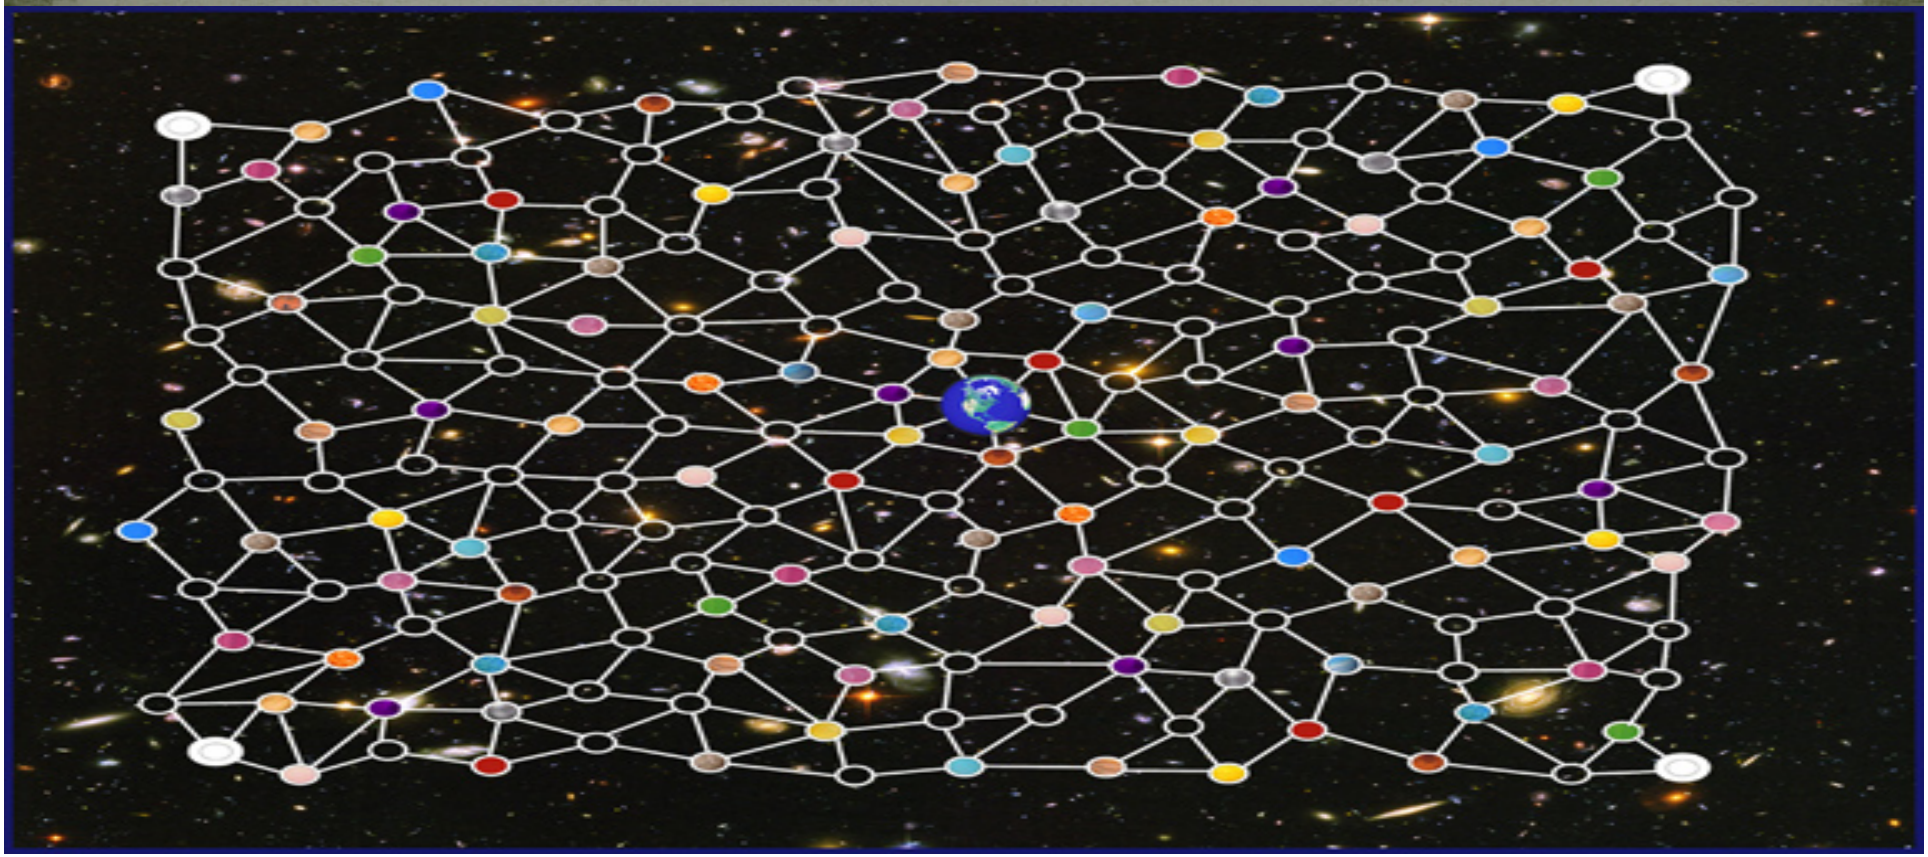

Applet Viewer: gui.GameApplet.class

Applet

Game Board

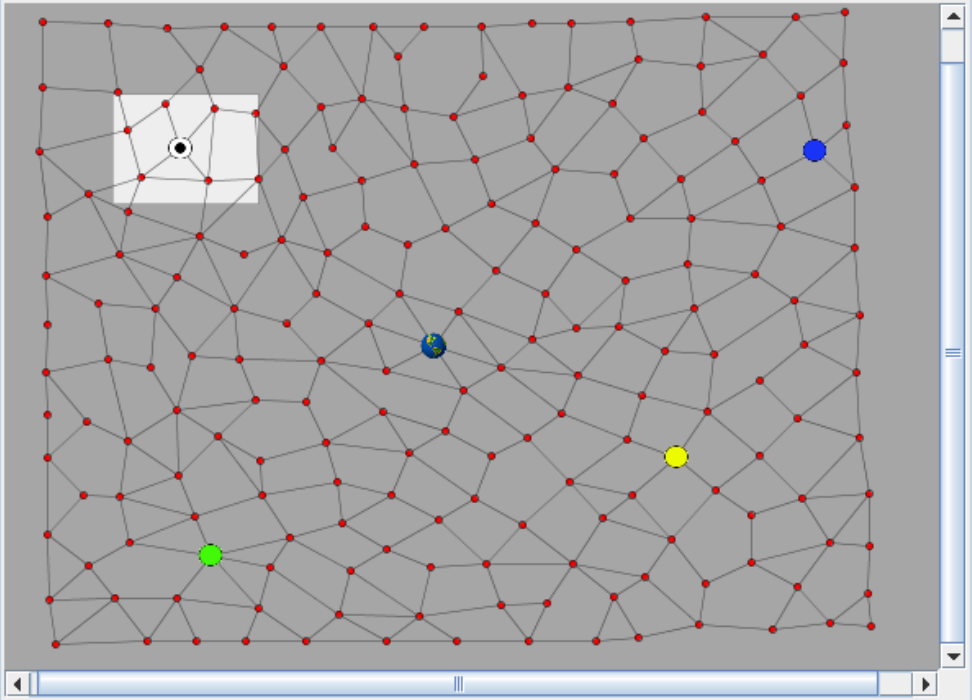

Resources

x 3

x 2

x 4

Notes

Player 2 won't share his resources.

Player 4 thinks Player 3 is the bad guy.

Trades

Trade With:

You Give:

0

0

0

You Get:

0

0

0

Reset

Offer Trade

Neighborhood

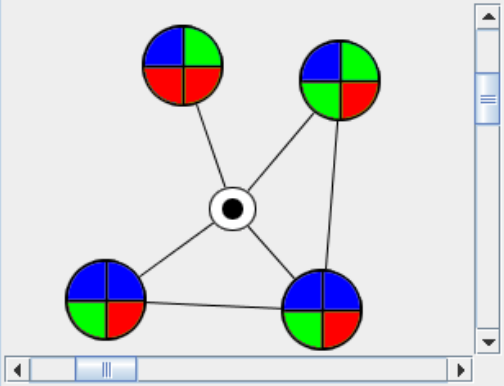

Chat

Player 1: Hi Everyone
Player 2: Hi, Good Luck!
Player 3: Can anyone give me a red resource?
Player 4: I can't, you might be the bad guy
Player 3: No I'm not! Maybe you are!
Player 1: I need a green

Type Here to Chat

Chat

Status Goes Here

Applet started.

Research Objectives and Results

- Understand how trust develops between people in social networks that change over time
- Build computational models of the dynamics of trust to guide decision-making and acting
- Paper in preparation for CHI-11

Trust in Social Groups with Hidden Roles/Objectives

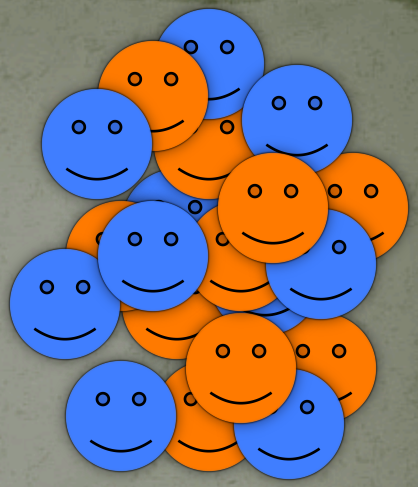

- Example:
  - Two groups of agents: Orange, Blue
    - All agents interact with each other for some number of iterations
    - $n_{Orange} + n_{Blue}$ agents; each agent only knows its role, but not others'
    - One can guess; but cannot be 100% sure until the end of the game
  - Winning condition
    - The group with high average accumulated payoff wins the game
- Research Question: how should a player represent and learn to adapt its trust to the others in the network?

Approach

- 2-population co-evolutionary model
  - "Gene" = initial trust value + trust-update rules
  - The fitness of an individual is computed by evaluating the individual against a number of individuals from the other group
- With this model, evolutionary computation will give us:
  - Automated generation/simulation of agent's trust models
  - Improve old models by arm race
  - Learning strategies for acting based on trust, from adversarial behavior
- We're currently working on the formalism of our model and algorithms, and are implementing the setup for pilot tests

Conclusions

- Described a new model for understanding cooperativeness in groups
  - Cooperativeness as a measure of intentions of the group
  - Other factors that go into intentions/objectives understanding?
- Ongoing: Learning how to adapt to an adversarial group
- Ongoing: Dynamics of social trust relationships among a group of agents
  - Among humans as well as mixed human/software agent groups
  - Evaluation framework: strategy games
- Future work: Develop a formalism (Behavior Network Diagrams) to analyze the above factors together
- Future Work: Using BNDs to do adversarial planning/decision making

Any Questions/Comments?

Motivating, abstract scenario

- Consider a humanitarian group trying to deliver supplies to a village that is suffering after a natural disaster.
- The transportation network provides many options for bringing in the support:
  - routes for helicopters, trucks, and people moving on foot.
- Several groups of people interacting
  - Relief people, Villagers, Militia, Others

Backup Slides

Importance of this work

- Described a formal model with insights from social and behavioral sciences
  - Analyzed the steady state behaviors
  - Showed the advantage of using social orientation of players as a metric for the society's cooperativeness level
    - While previous works using average payoff as a metric of the degree of cooperation may not be accurate sometimes
- We are in the process of performing more extensive evaluations of this controversial result
- However, this result suggests that understanding, game-theoretically, the dynamics of a group's behavior has some open, interesting issues
  - Next slides on two ongoing works

An example: a Fair player in IPD

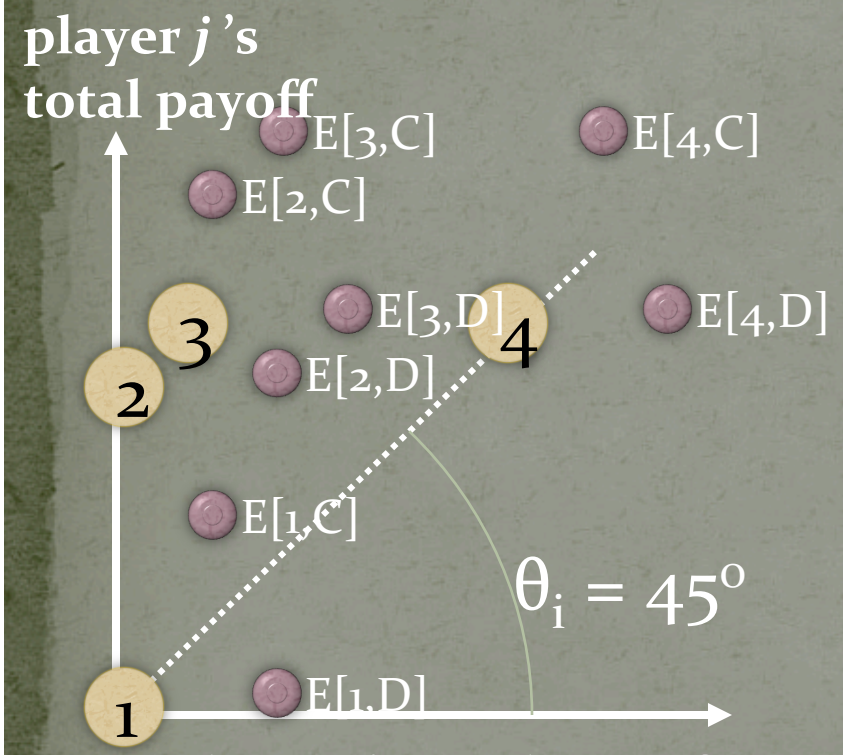

- Player i is fair, i.e., $\theta_i = 45^o$
- Player j is random
- $g_1 = \langle 0, 0 \rangle$
- i chooses C (j chooses D)
- $g_2 = g_1 + \langle 0, 5 \rangle = \langle 0, 5 \rangle$
- i chooses D (j chooses D)
- $g_3 = g_2 + \langle 1, 1 \rangle = \langle 1, 6 \rangle$
- i chooses D (j chooses C)
- $g_4 = g_3 + \langle 5, 0 \rangle = \langle 6, 6 \rangle$
- i chooses C
- Similar to Tit-for-Tat

What happens in Steady State?

- Our theoretical analyses showed predictions on the behavior conditioned on initial social orientations

2x2 symmetric game		Player 2	
		Cooperate (C)	Defect (D)
Player 1	Cooperate (C)	(R, R)	(S, T)
	Defect (D)	(T, S)	(P, P)

- Five possible cases for each game in steady state:
  - If both i & j are pro-self, then both players always defect
  - If both i & j are pro-social, then both always cooperate
  - If i is pro-social, but j is Pro-self, then Player i gets $r_i(T+S)$, and Player j gets $(1-r_i)(T+S)$
  -

Cycle of Cooperation and Defection

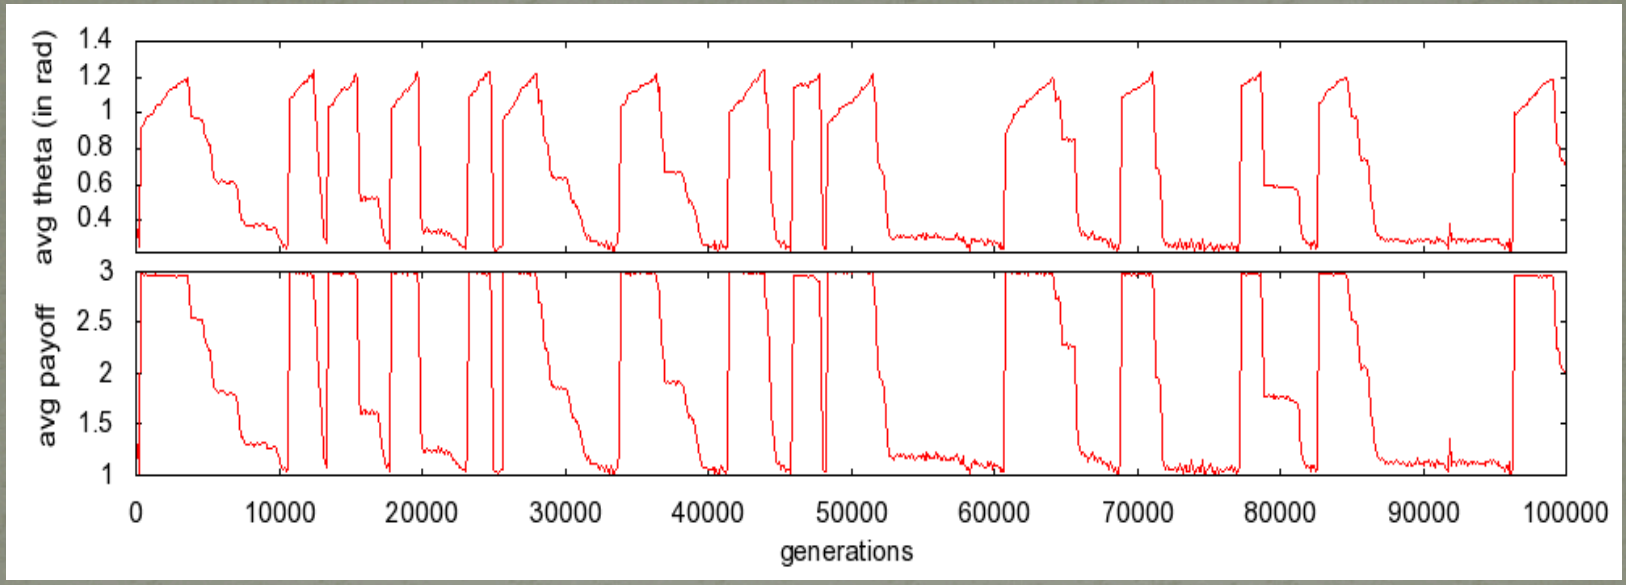

- Results confirm the previous findings in Evolution of Cooperation (Axelrod 1980):
  - Average payoff varies between P and R, which correspond to full defection and full cooperation, respectively

Example

- Each agent keeps track of trustworthiness, t_i, of other agents, $-1 \leq t_i \leq 1$
  - Agent 2 has $t_1 = 1$ means 2 has high trust in agent 1
    - i.e., agent j belief that agent i's behavior is similar to its, so they're in the same team
  - Agent 2 has $t_1 = -1$ means 2 has low trust in 1
    - i.e., agent j belief that agent i's behavior will against agent j's plans
- After each iteration, t_i is updated based on the history, according to the following updating rule
  - H -> Δt_H, where H is the history of interaction
  - E.g., t_2=0.1 and ⟨(C,C),…,(C,C) ⟩-> 0.1, then t_2 is updated to 0.2 if the history is ⟨(C,C),…,(C,C) ⟩

Related work on modeling agent's behavior

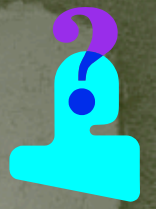

- Game-theoretical approaches
  - Use rationality assumption, i.e., maximizing their own utility
- Several models for evolution of cooperation
  - Examples
    - Lookup table e.g., [Axelrod 87]
    - Probabilities of cooperation e.g., [Nowak and Sigmund 93]
    - Finite state automata e.g., [Linster 92]
    - Neural networks e.g., [Harrald and Fogel 96]
  - Most of them did not explicitly take the payoff of the other player into account
- Research in social and behavioral sciences
  - E.g., Social Value Orientation (SVO) Theory

Game Play:

- There are three good players and one bad player.
- All players are heading to the earth at the center.
  - The bad player wins if he gets to the center ahead of at least one good player.
  - Otherwise the good team wins.
- Players draw cards at the beginning of the game that will state if they are good or bad.
  - They keep these cards private so no one truly knows the goodness or badness of other players.
- At the beginning of each round, the players each roll the die.
  - The player draws the number of resource cards as is shown on the die.
  - All players roll for resources in all rounds, even once they have reached the center

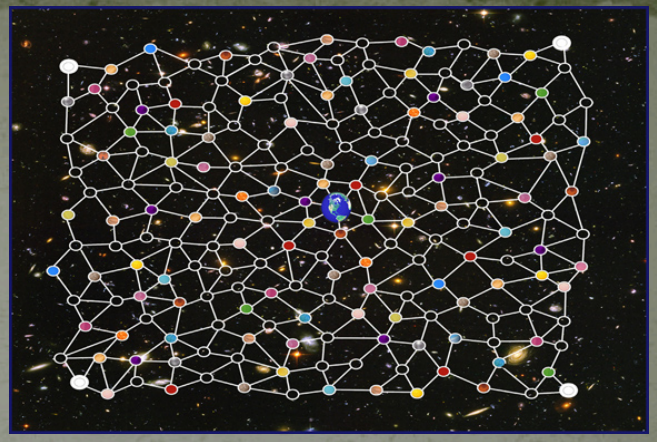

- After rolling, players draw requirement cards for each path they might want to cross.
- To move, a player must turn in resource cards to pay the cost of the path they will take, and they must show the requirement card at the same time. All cards go into a discard pile.
- Resources can be exchanged between players.
  - They can be traded, given away, or exchanged with promise of something in return later.
  - There are no rules to prevent players from lying or going back on promises - anything is fair!

Trust-Based Applications and Interfaces for the Social Web

John O'Donovan,
Dept. of Computer Science. University of California, Santa Barbara.

Workshop on Computing with Social Trust, University of Maryland, College Park. July 29th 2010

Outline

“We need better ways to model and represent trust…”

Past Work…

(At University College Dublin, and Univ. Southern California)

Trust on the Social Web (PhD Thesis)

Example Studies: Trust Modeling for Recommender Systems

Trust Modeling for Online Auctions

Currently…

(At Univ. of California Santa Barbara)

Trust-based Interfaces for Recommender Systems

Interactive analysis of trust networks.

Example: SmallWorlds

Next…

Transparent Trust Models?

Portable and Standardized?

The role of the user interface for trust and uncertainty information.

Linked Data & semantic web solutions?

Examples: WiGis and WiGiPedia.

Our Knowledge Repository- 1910

WHEN IN DOUBT–"LOOK IT UP" IN

The

Encyclopaedia Britannica

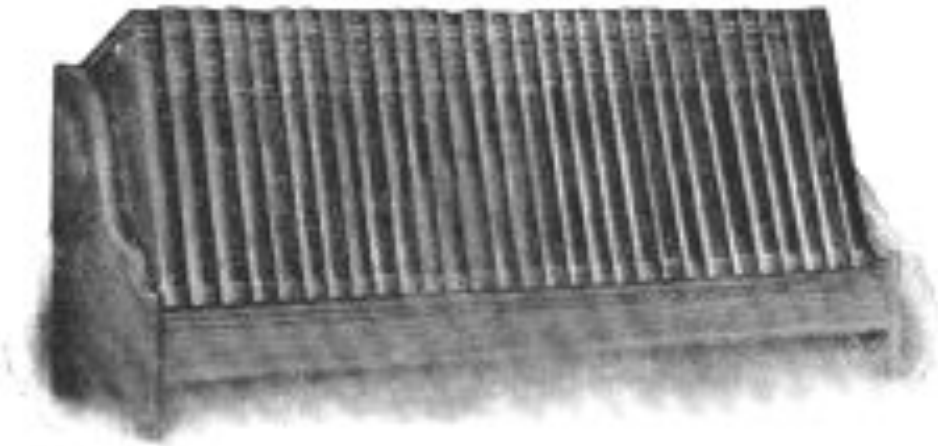

(New 11th Edition) issued 1910-11 by the
CAMBRIDGE UNIVERSITY PRESS (England)

The Sum of Human Knowledge

29 volumes, 28,150 pages, 44,000,000 words of text. Printed on thin, but strong opaque India paper, each volume but one inch in thickness.

THE BOOK TO ASK QUESTIONS OF

FOR READING OR FOR STUDY

The "Traditional" Web (aka 1.0)

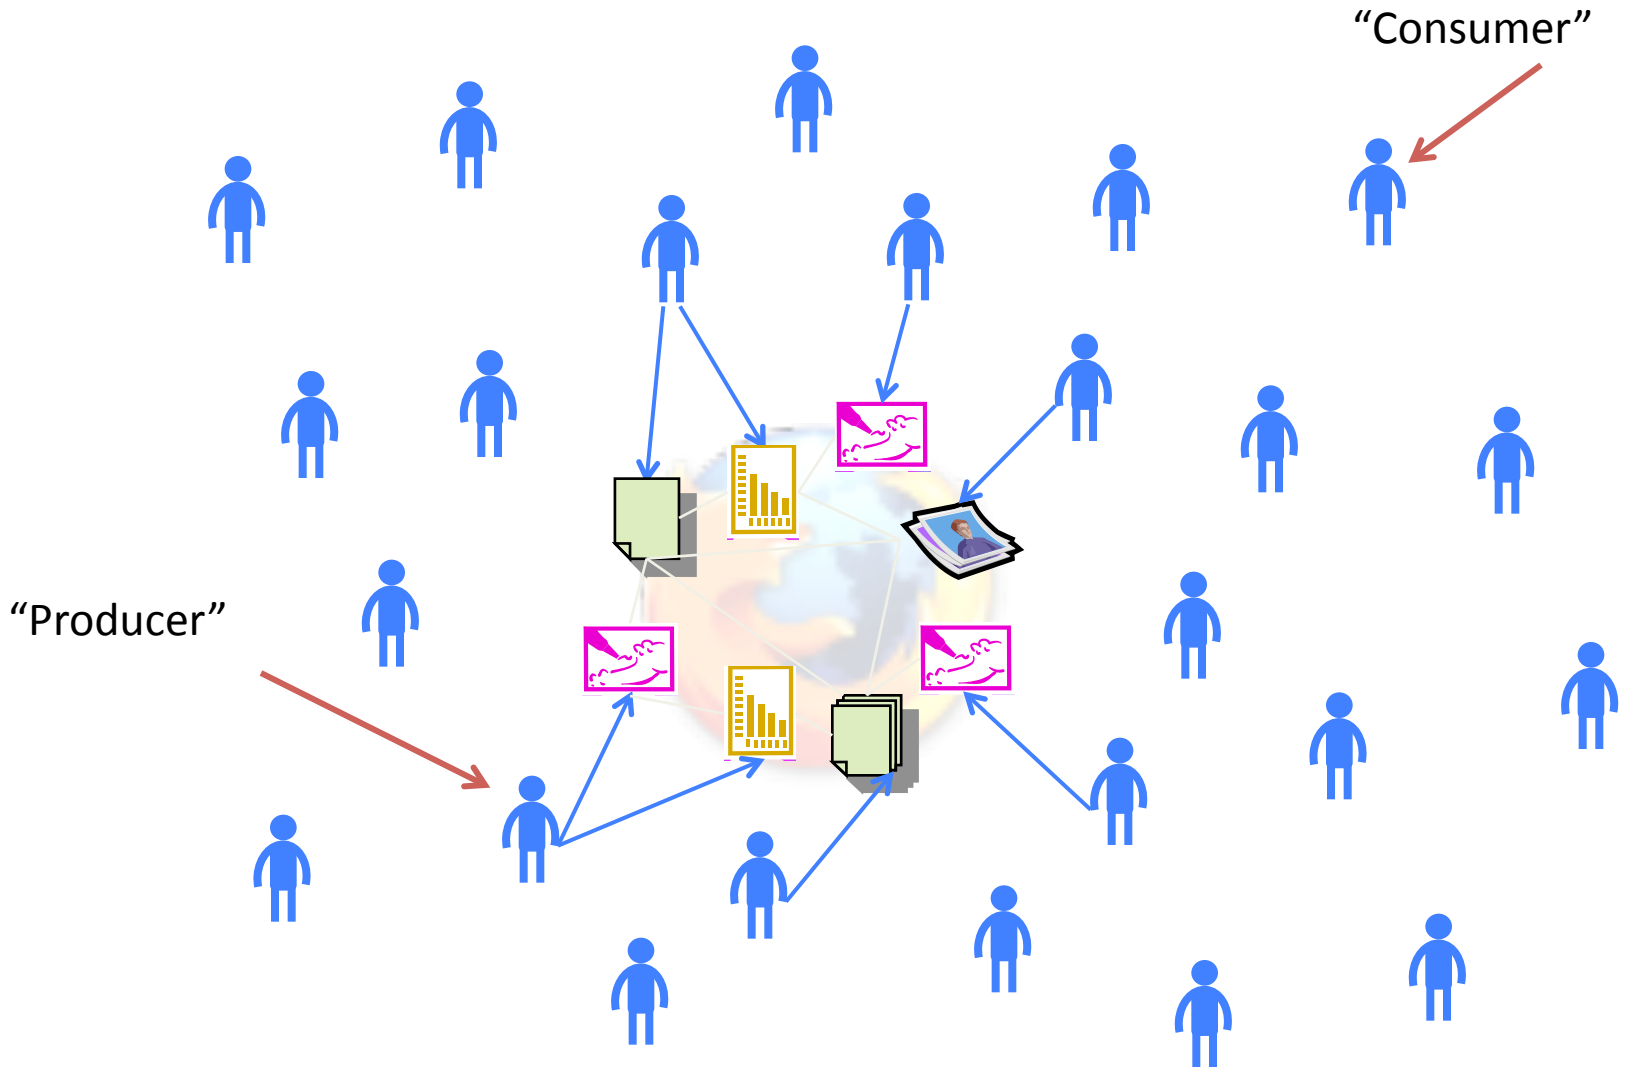

The "Social" Web (aka 2.0)

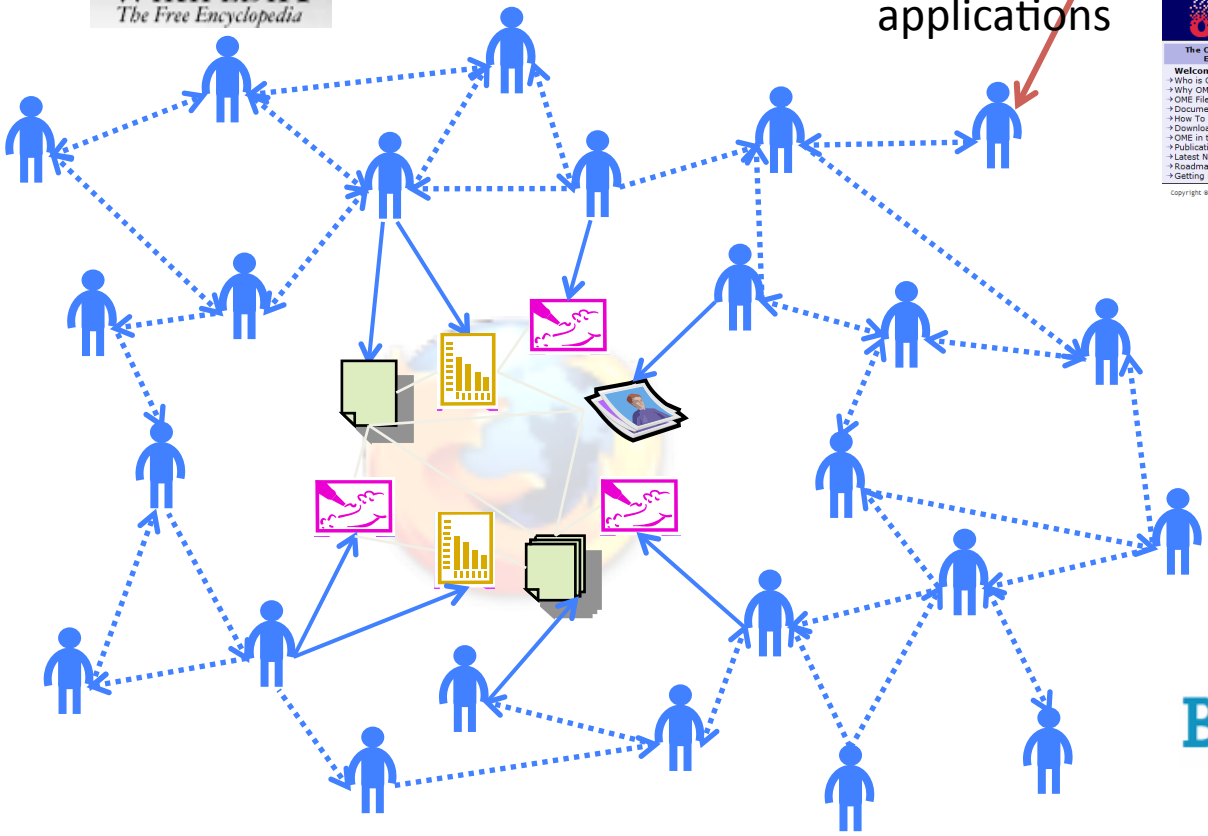

Specialist applications

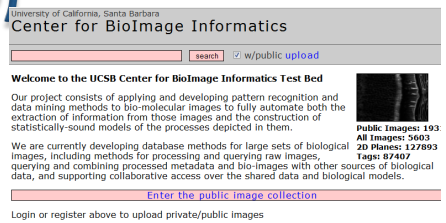

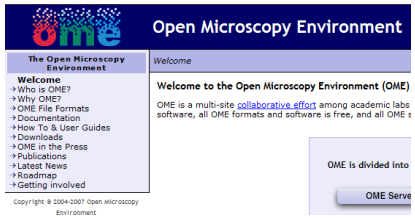

flickr

WIKIPEDIA
The Free Encyclopedia

YAHOO! ANSWERS

vimeo

digg

Bloglines

Technorati

YouTube
Broadcast Yourself™

twitter

Google's PageRank:

"Trust" of web *pages* providing info

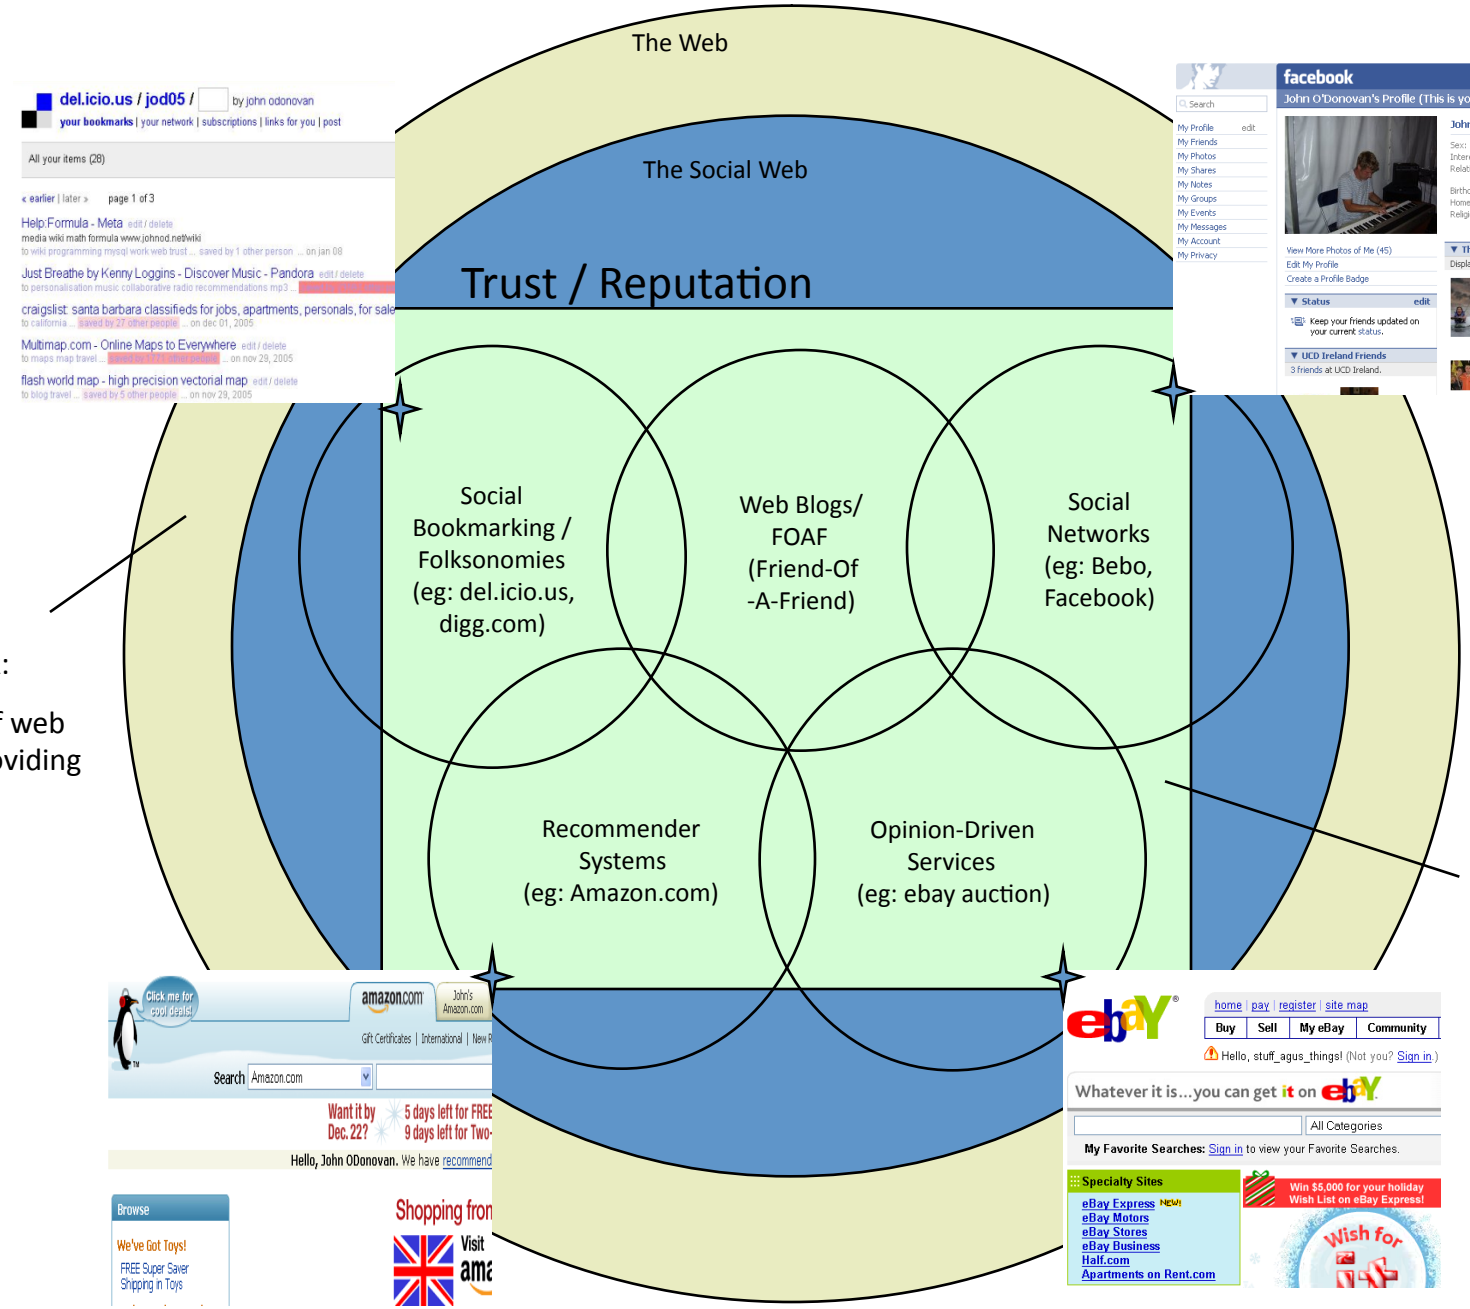

Social Web: 'Trust' of *people* contributing info

Previous Work on Trust.

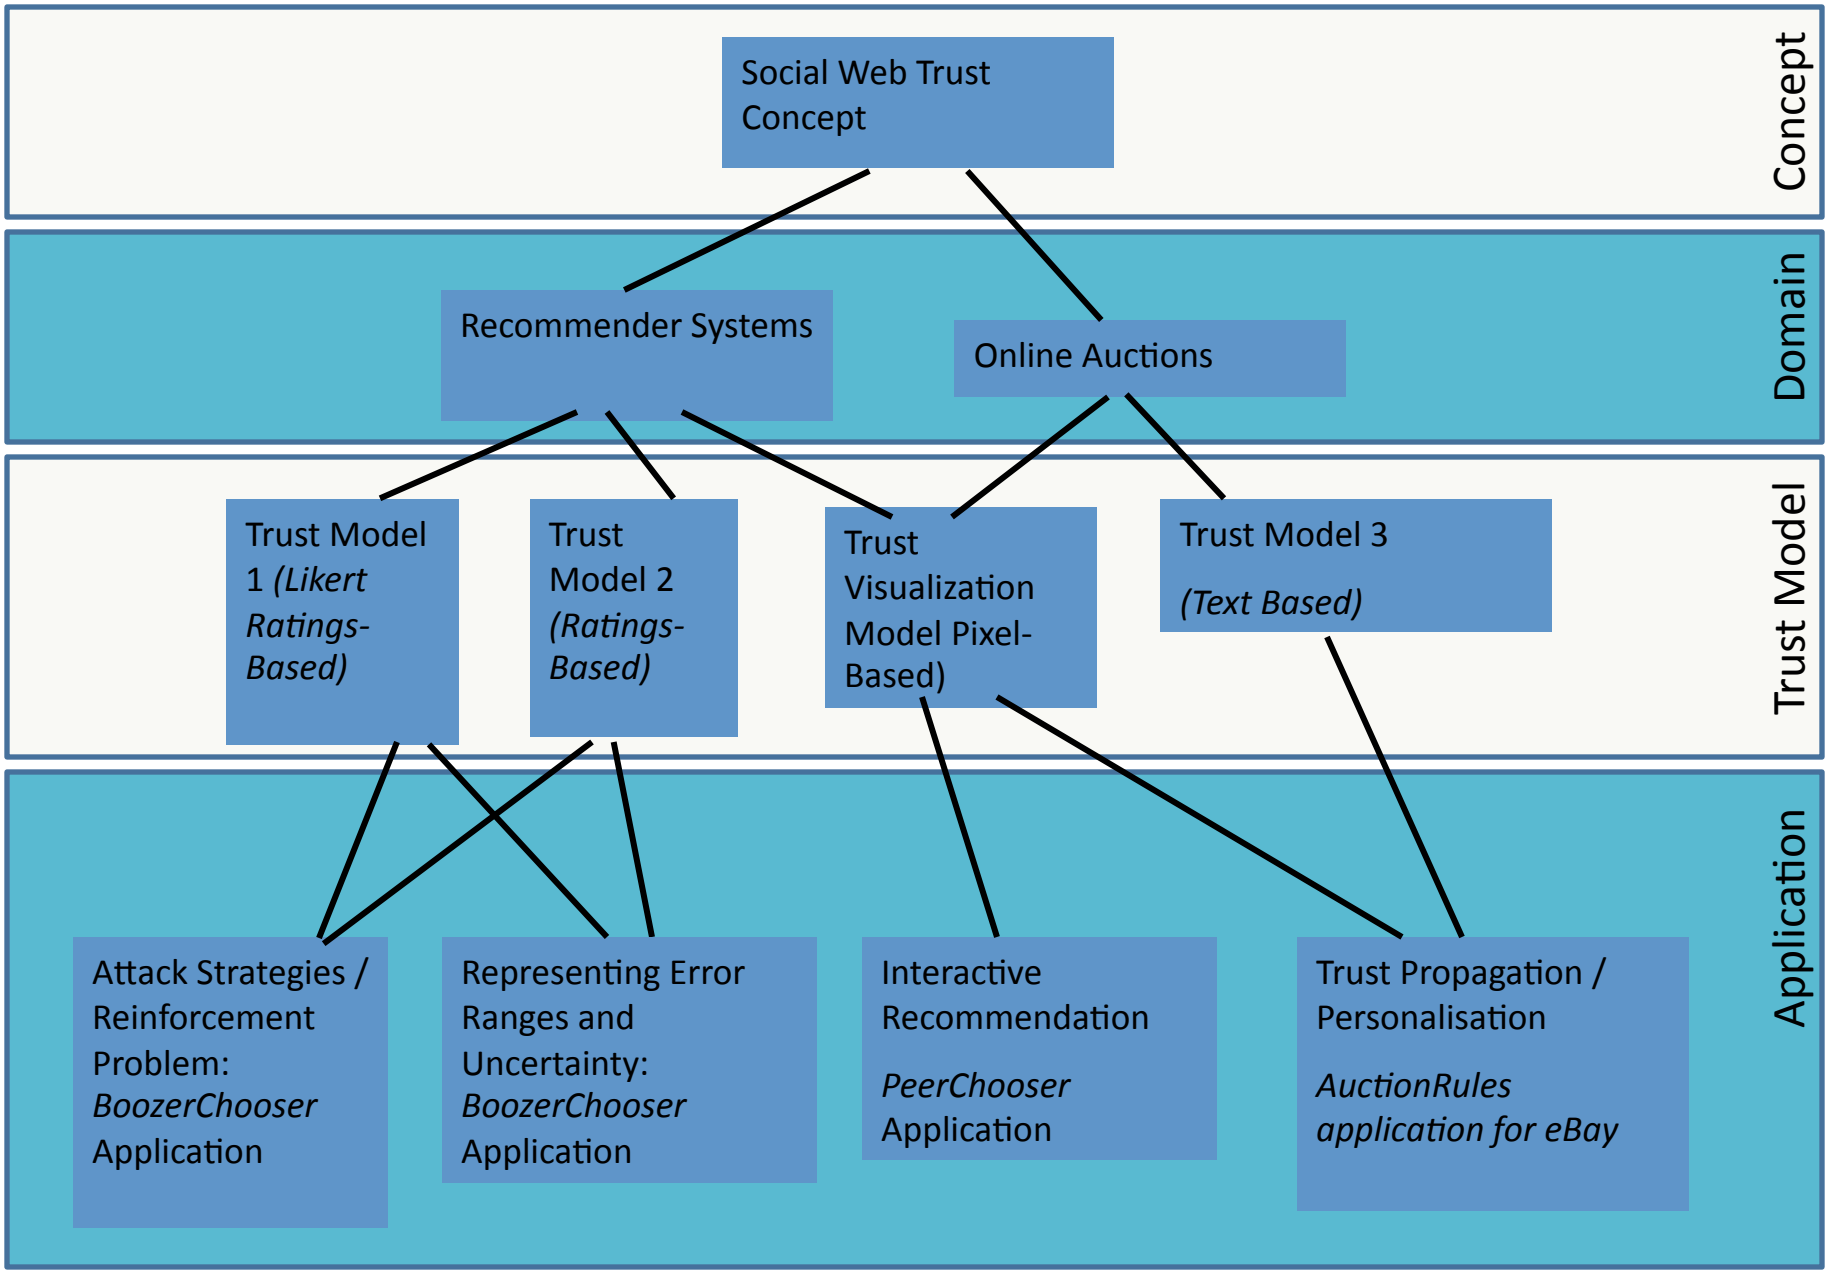

Example 1: Modeling trust in a recommender system.

•Collaborative filtering algorithms rely too heavily on profile similarity.

•Trustworthiness of users is an important factor to be considered.

•Trust can be non-invasively inferred from user rating histories.

•Trust values can easily be incorporated into the mechanics of a standard collaborative filtering algorithm to yield better robustness and predictive accuracy results.

Standard Collaborative Filtering

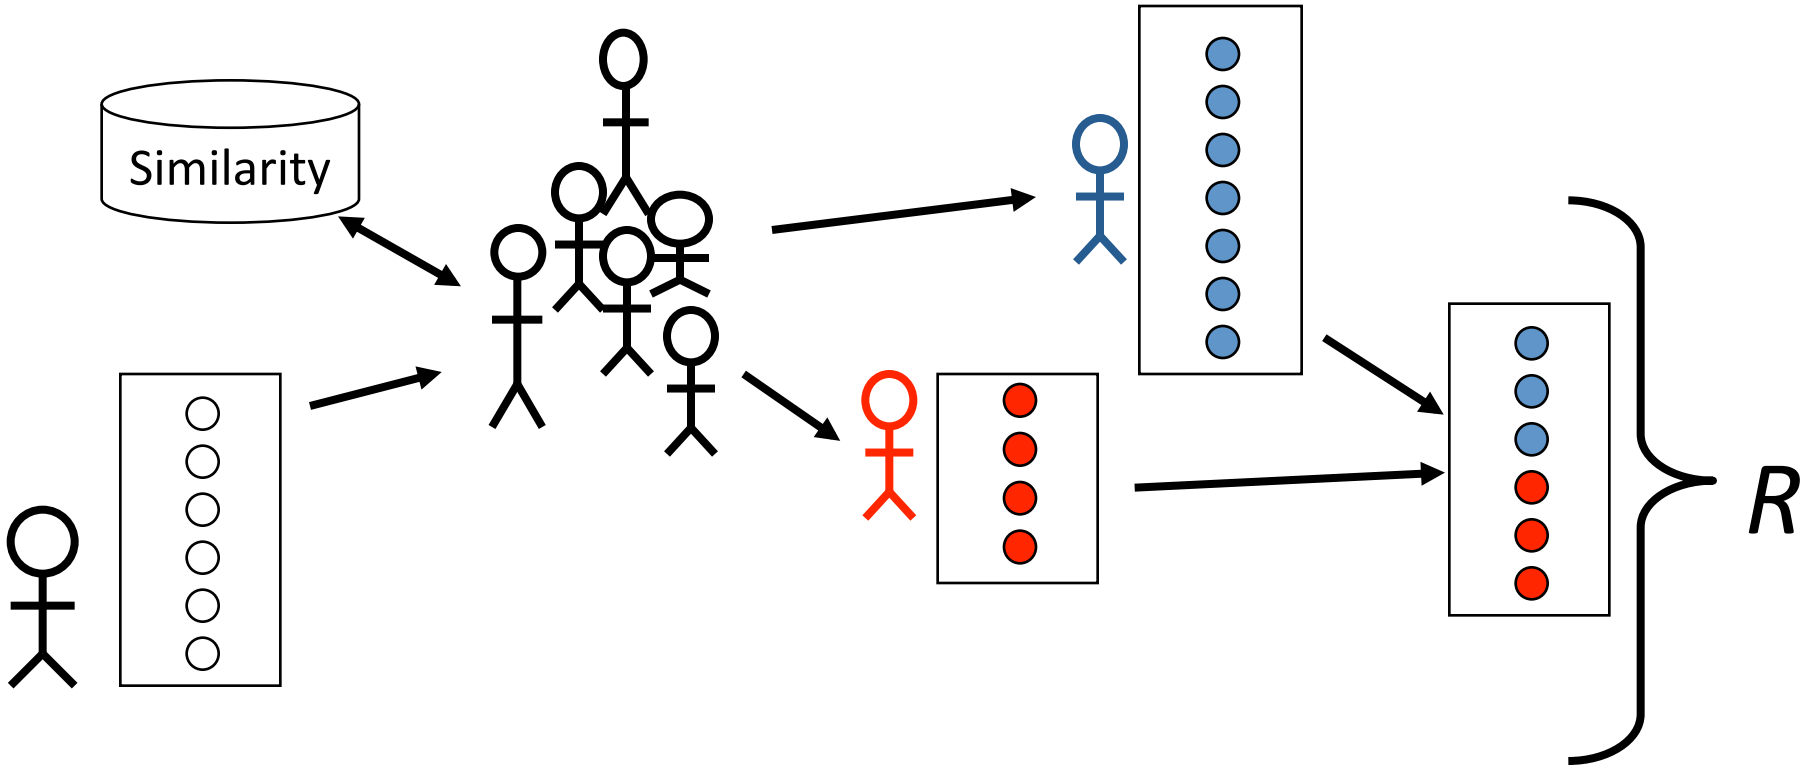

Trust-Based Collaborative Filtering

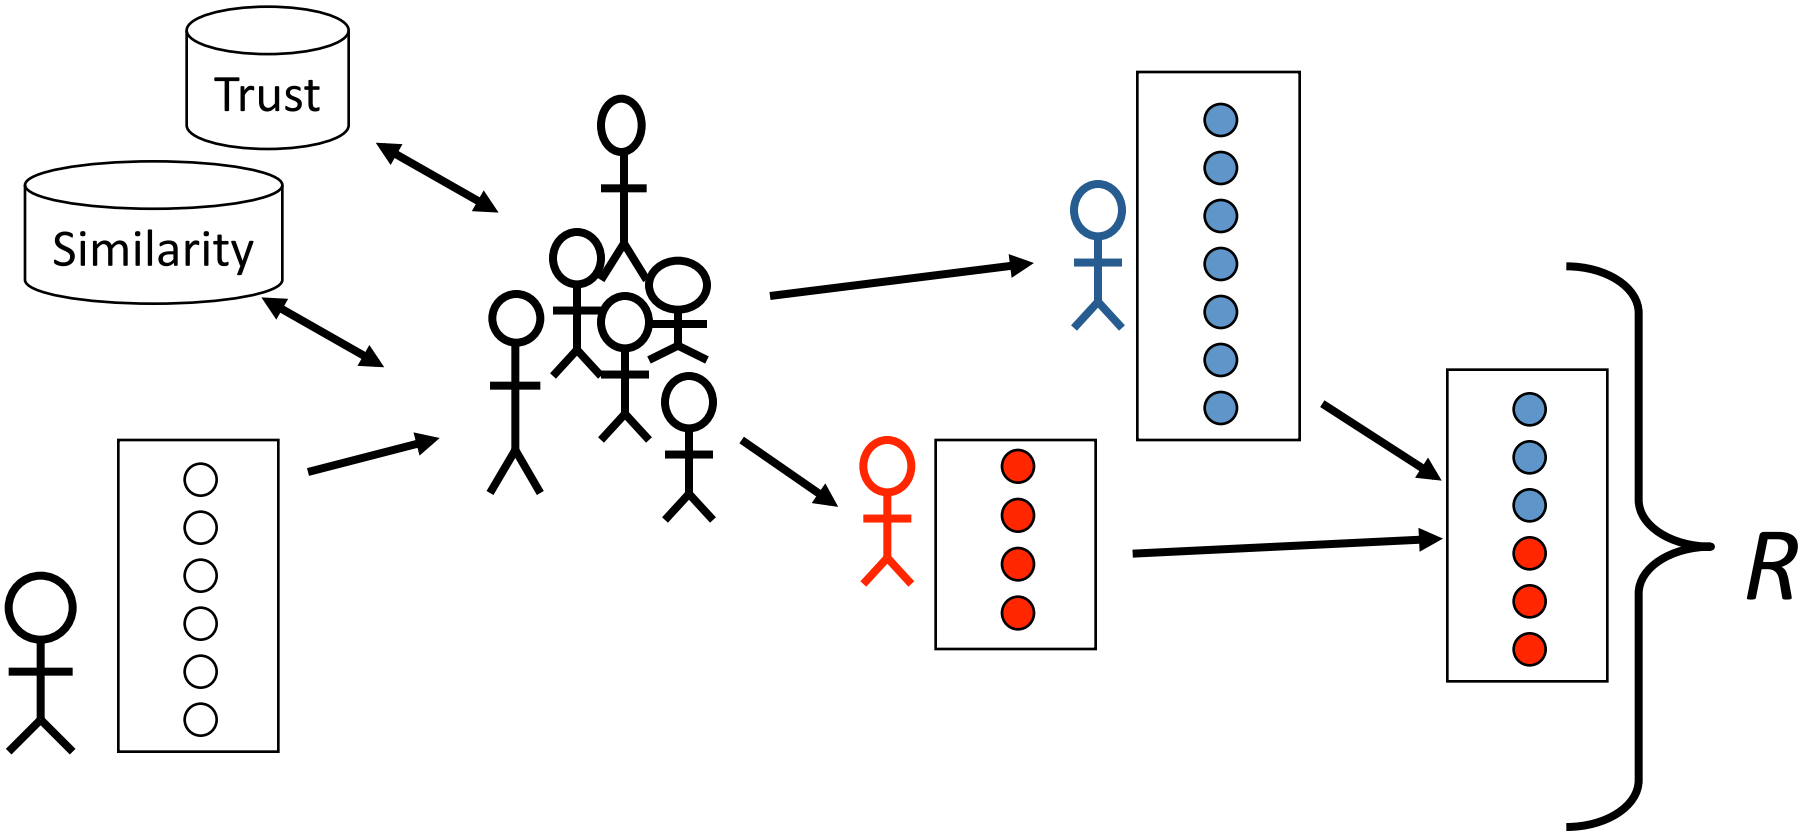

Building the trust model

$$Trust^{I}(p,i) = \frac{|\{(c_k, i_k) \in CorrectSet(p) : i_k = i\}|}{|\{(c_k, i_k) \in RecSet(p) : i_k = i\}|}$$

$$Trust^{P}(p) = \frac{|CorrectSet(p)|}{|RecSet(p)|}$$

Trust-Based Collaborative Filtering

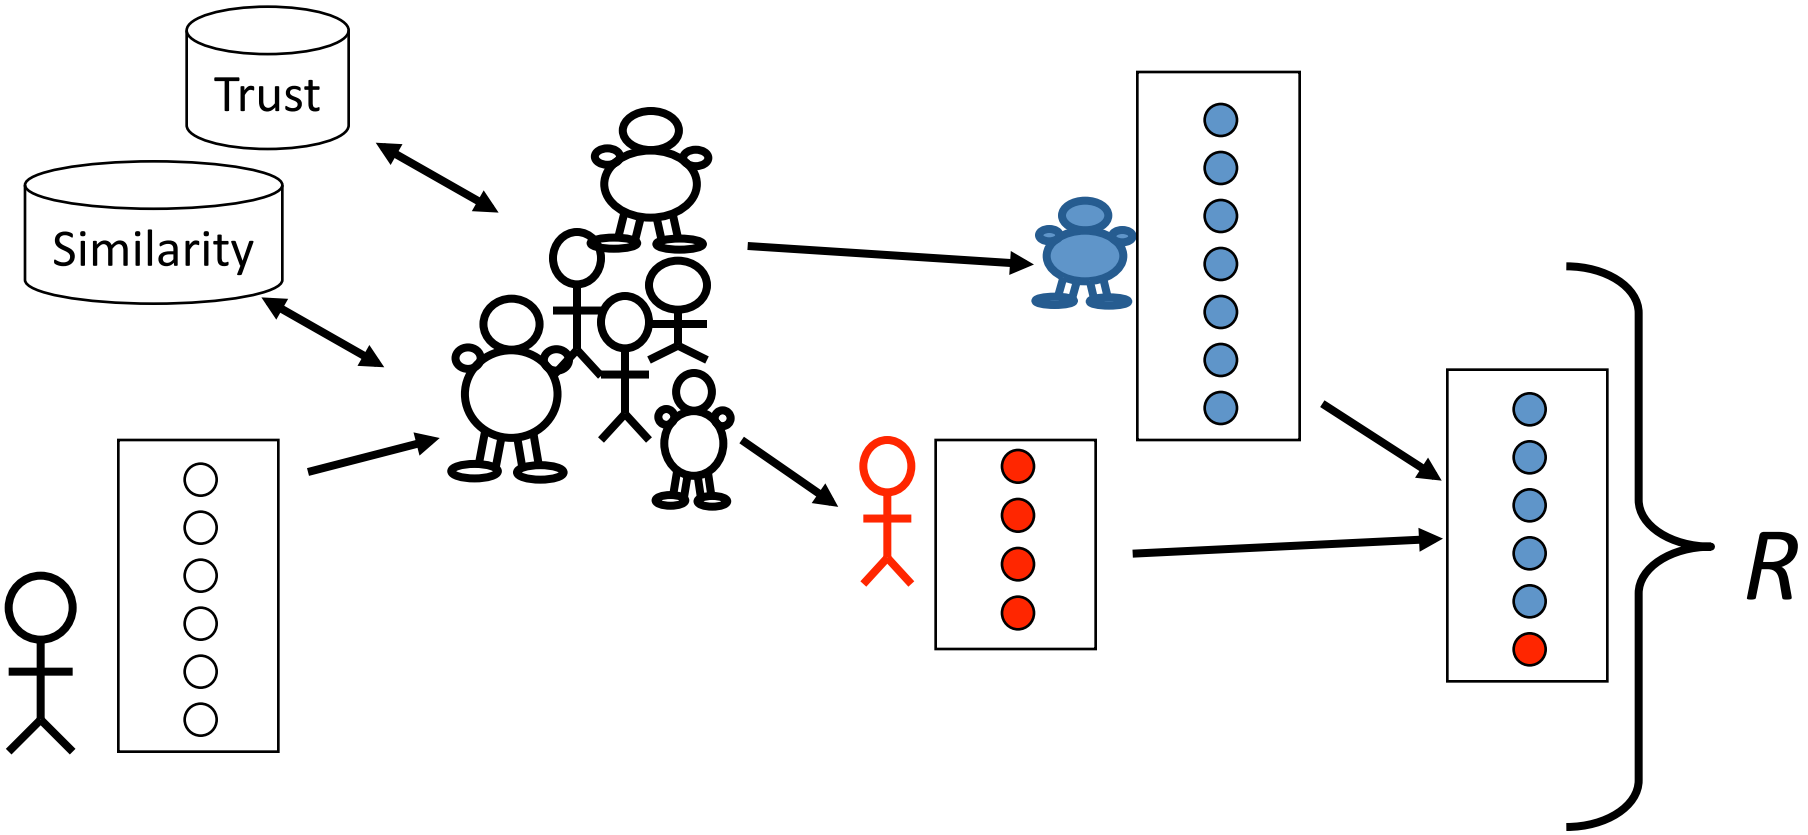

Example 2: AuctionRules Trust-mining Algorithm

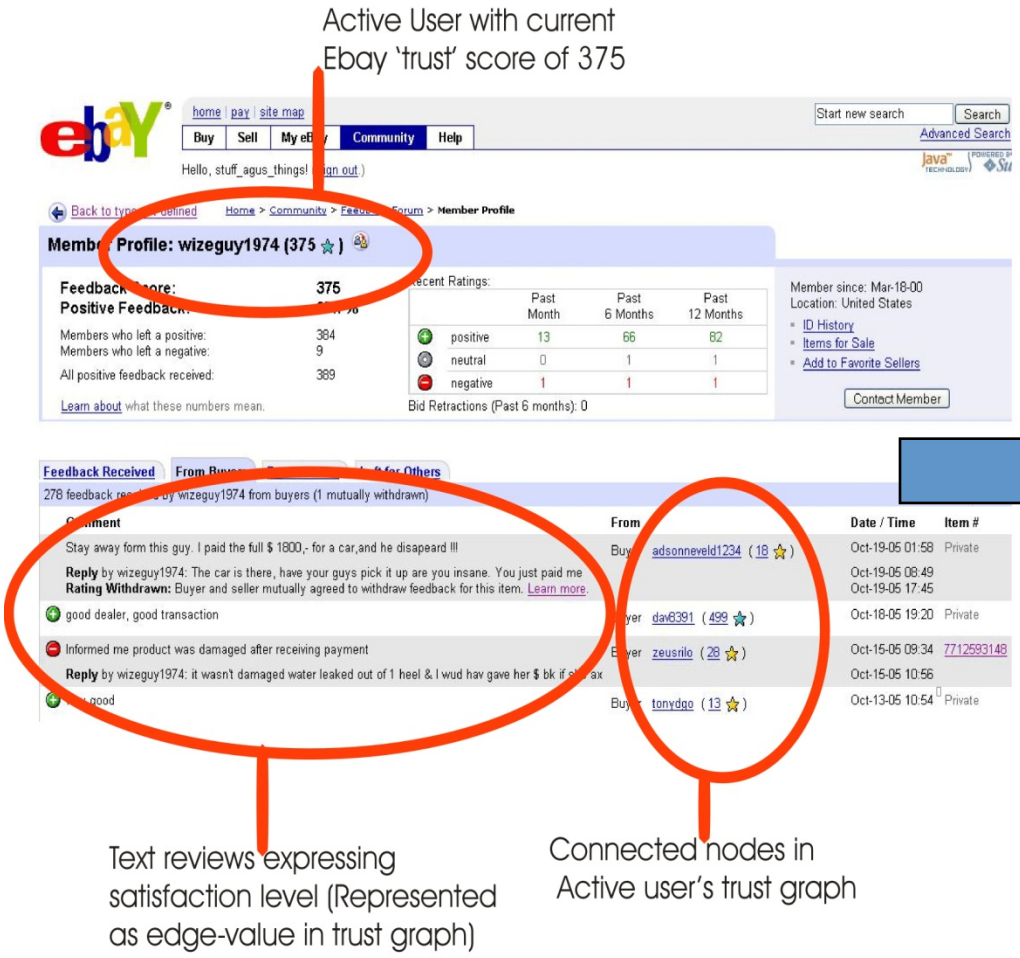

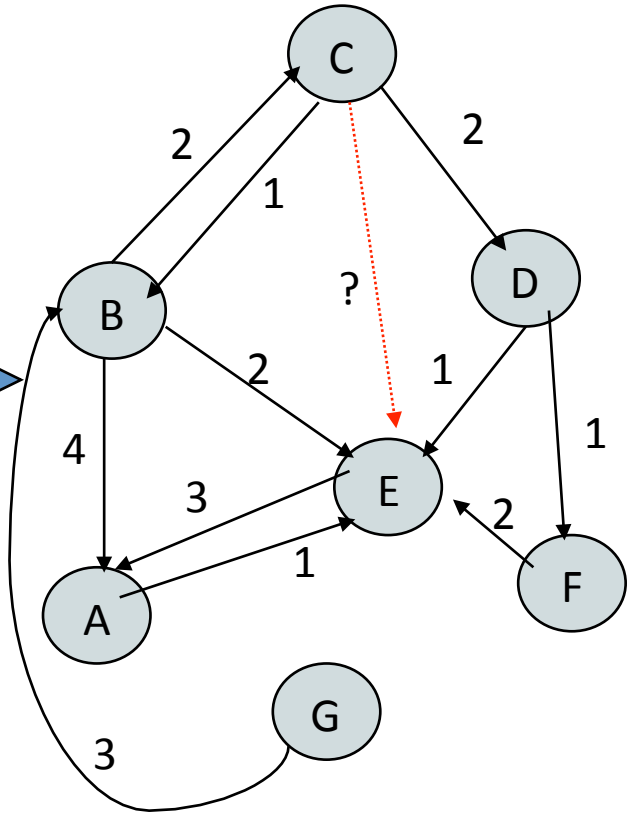

Example 2: AuctionRules Trust-mining Algorithm

We designed a lightweight NLP classification algorithm specifically for extracting the polarity of sentiment in online auction comments on eBay…

Representing the Auction as a Set of Features

1. *Item* - The quality/condition of the product being bought or sold. (*item, product*)
2. *Person* - The person the user makes the transaction with. (*buyer, seller, eBayer*, *dealer*)
3. *Cost* - Cost of item, cost of shipping, hidden costs etc. (*expense, cost*)
4. *Shipping* - Delivery of the item, security, time etc. (*delivery, shipping*)
5. *Response* - Communication with the other party, emails, feedback comment responses. (*response, comment, email, communication*)
6. *Packaging* - The packaging quality/condition of the item *packaging*
7. *Payment* - how the payment will be made to the seller, or back to buyer for return (*payment*)
8. *Transaction* - the overall transaction quality (*Service, Transaction, Business*)

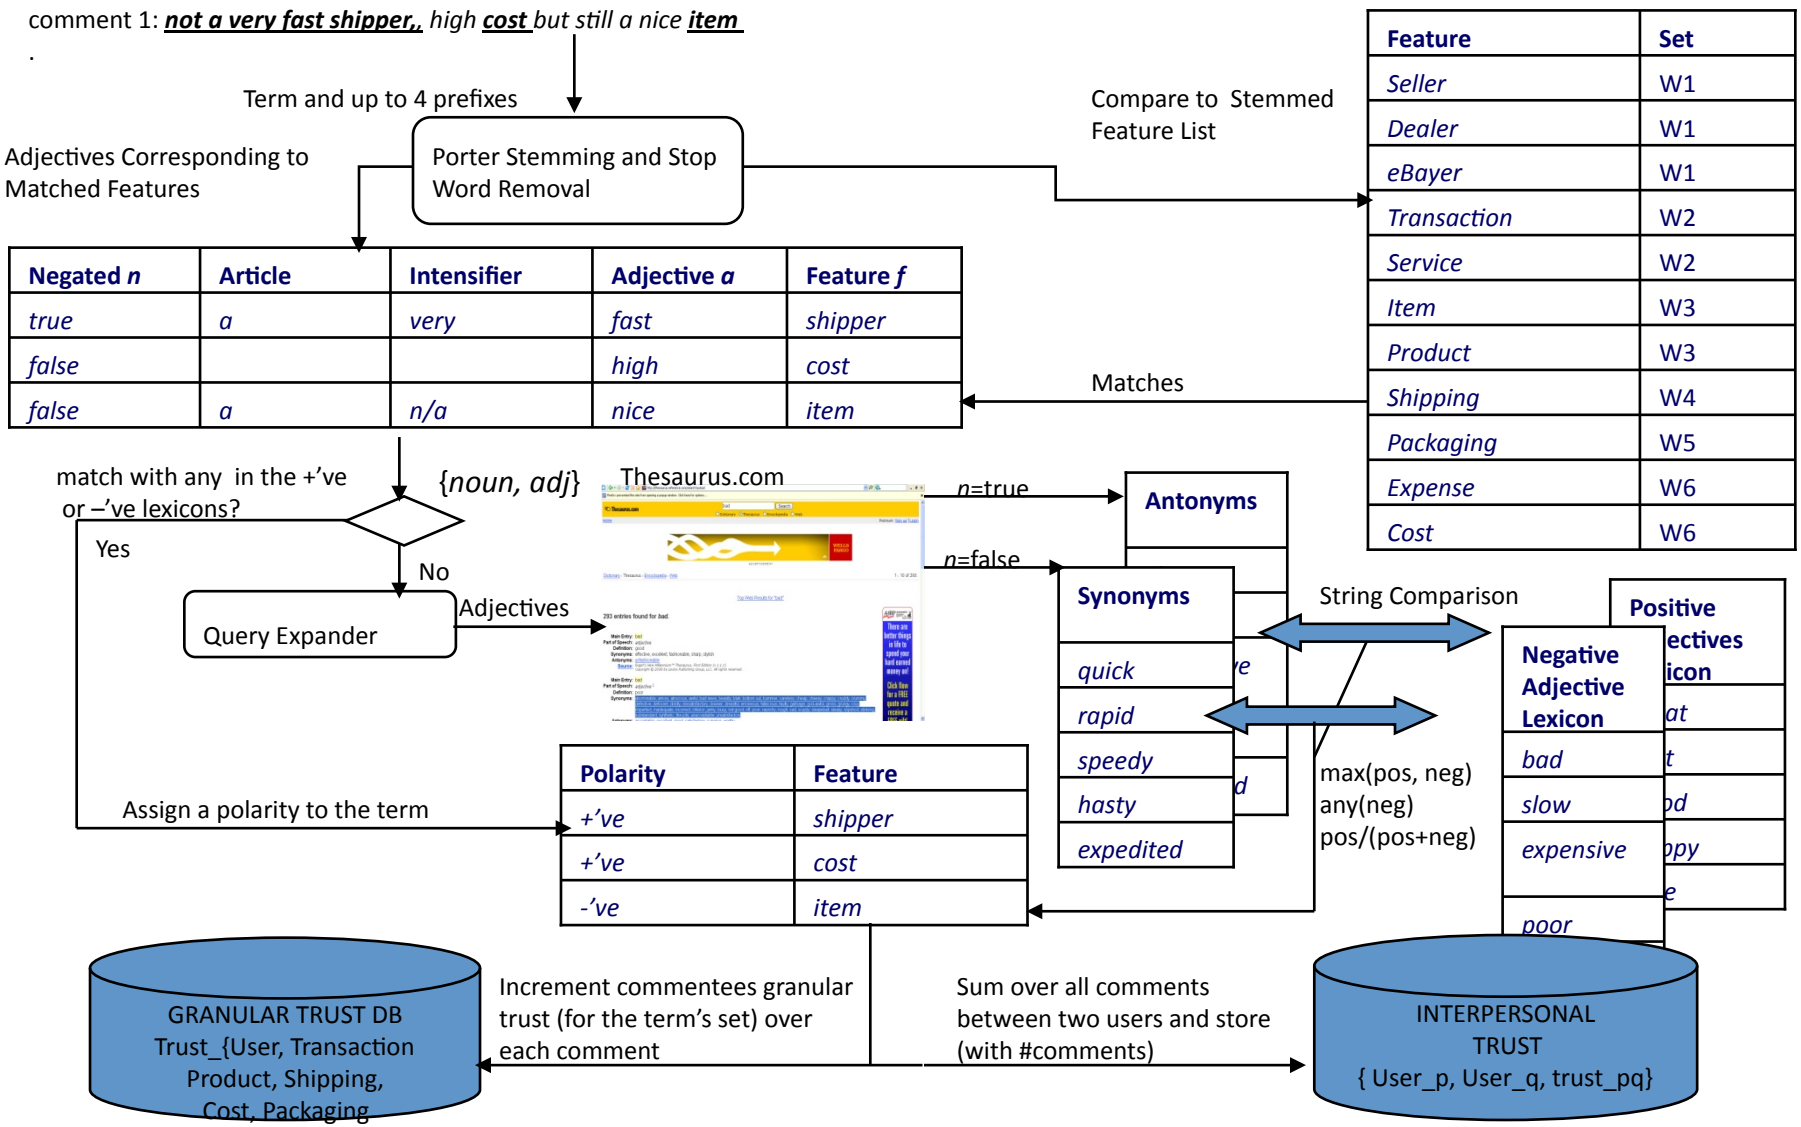
comment 1: not a very fast shipper,, high cost but still a nice item
.
Term and up to 4 prefixes
Porter Stemming and Stop Word Removal
Adjectives Corresponding to Matched Features
Compare to Stemmed Feature List
Feature
Set
Seller
W1
Dealer
W1
eBayer
W1
Transaction
W2
Service
W2
Item
W3
Product
W3
Shipping
W4
Packaging
W5
Expense
W6
Cost
W6
Negated n
Article
Intensifier
Adjective a
Feature f
true
a
very
fast
shipper
false
high
cost
false
a
n/a
nice
item
Matches
match with any in the +'ve or –'ve lexicons?
{noun, adj}
Thesaurus.com
n=true
n=false
Antonyms
Synonyms
quick
rapid
speedy
hasty
expedited
String Comparison
Yes
No
Query Expander
Adjectives
Positive Adjectives Lexicon
Negative Adjective Lexicon
bad
slow
expensive
poor
max(pos, neg)
any(neg)
pos/(pos+neg)
Assign a polarity to the term
Polarity
Feature
+'ve
shipper
+'ve
cost
-'ve
item
GRANULAR TRUST DB
Trust_{User, Transaction Product, Shipping, Cost, Packaging
Increment commentees granular trust (for the term's set) over each comment
Sum over all comments between two users and store (with #comments)
INTERPERSONAL TRUST
{ User_p, User_q, trust_pq}

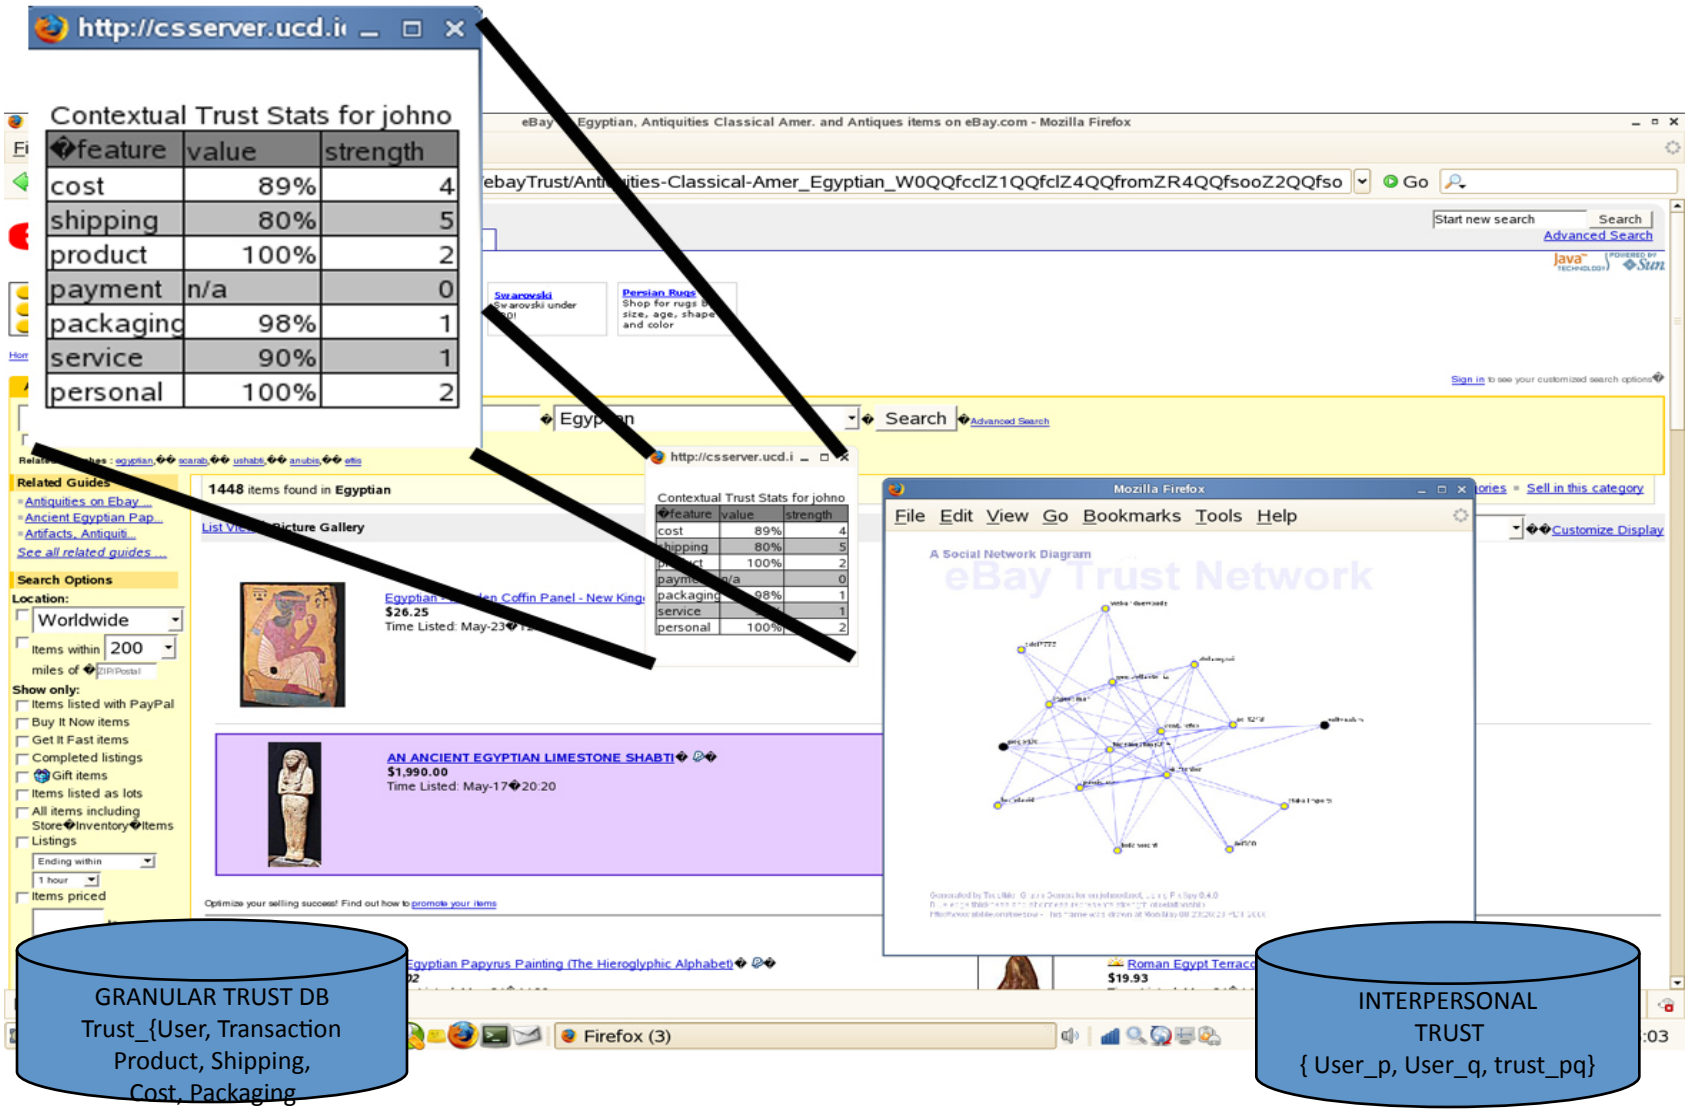
http://csserver.ucd.ie
Contextual Trust Stats for johno
feature | value | strength
cost | 89% | 4
shipping | 80% | 5
product | 100% | 2
payment | n/a | 0
packaging | 98% | 1
service | 90% | 1
personal | 100% | 2
eBay Egyptian, Antiquities Classical Amer. and Antiques items on eBay.com - Mozilla Firefox
Mozilla Firefox
File Edit View Go Bookmarks Tools Help
A Social Network Diagram
eBay Trust Network
1448 items found in Egyptian
AN ANCIENT EGYPTIAN LIMESTONE SHABTI
$1,990.00
Time Listed: May-17 20:20
Egyptian Papyrus Painting (The Hieroglyphic Alphabet)
Roman Egypt Terrac
$19.93
Firefox (3)
GRANULAR TRUST DB
Trust_{User, Transaction
Product, Shipping,
Cost, Packaging
INTERPERSONAL
TRUST
{ User_p, User_q, trust_pq}

Current Work

at CS Dept, UCSB

- Broad area of interactive network analysis (with trust data) for social web and other applications.

Five Ongoing Projects:

- WiGi: Web-based Interactive Graph Interfaces
- WiGiPedia: Interactive visualizations for eliciting trustworthy semantic data in WikiPedia.
- SmallWorlds: Interactive Trust-based recommender system for Facebook.
- TopicNets: Visual, Interactive and Real-Time Exploration of Topic-based relations in diverse datasets.
- Provenance Visualization: Interactive visualization of provenance and trust data for complex scientific process flows

WiGis Framework:

Scalable Interactive Visualization of >1M networked entities
An order of magnitude more scalable than the next best web based graph tool

Customizable Semantic Framework. Map any data field- (eg: Trust, Similarity) to a Graph Dimension

Rich search functionalities, Node, Edge, Group and Shortest Path Highlighting

Native in all major browsers, with no plugins (flash, java etc)

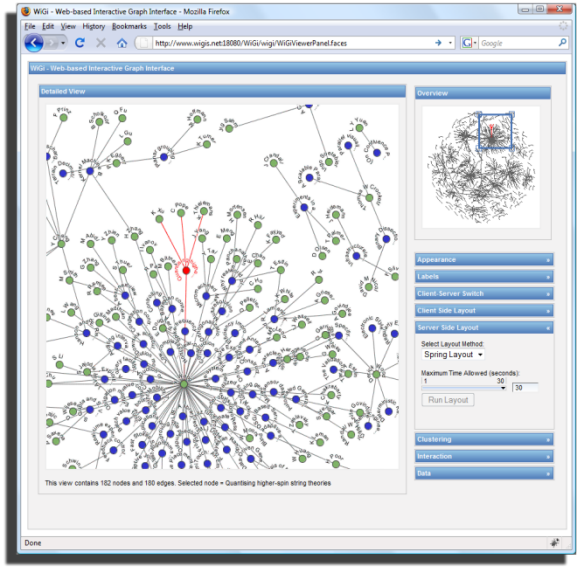

Overview and Zoom Navigation

Visualize Remote algorithms running on local data in real time

Feature Rich: Layout, Clustering and Interaction algorithms

Seamless Transition between client (local) and server based (remote) data models. Graph representations synchronized through AJAX.

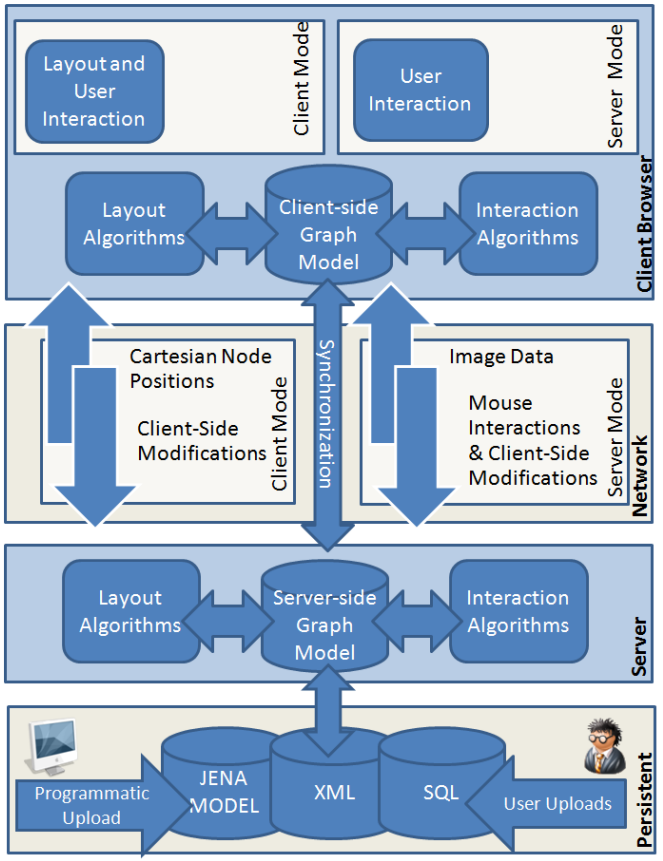

WiGipedia

article | discussion | edit this page | history

Try Beta | Log in / create account

Pink Floyd

From Wikipedia, the free encyclopedia

This article is about the rock band. For the fictional character, see Pink Floyd (fictional character).

Pink Floyd were an English rock band who earned recognition for their psychedelic music in the late 1960s, and as they evolved in the 1970s, for their progressive rock music. Pink Floyd's work is marked by the use of philosophical lyrics, sonic experimentation, innovative album cover art, and elaborate live shows. One of rock music's most critically acclaimed and commercially successful acts, the group has sold over 200 million albums worldwide, including 74.5 million certified units in the United States.

Pink Floyd were formed in 1965, and originally consisted of university students Roger Waters, Nick Mason, Richard Wright, and Syd Barrett. The group were a popular fixture on London's underground music scene, and under Barrett's leadership released two charting singles, "Arnold Layne" and "See Emily Play", and a commercially and critically successful debut album, *The Piper at the Gates of Dawn*. In 1968, guitarist and singer David Gilmour joined the line-up, and Barrett was removed due to his increasingly erratic behaviour. Following Barrett's departure, bass player and singer Roger Waters became the lyricist and dominant figure in the band, which thereafter achieved worldwide critical and commercial success with the concept albums *The Dark Side of the Moon*, *Wish You Were Here*, *Animals*, and rock opera *The Wall*.

Wright left the band in 1979, and Waters in 1985, but Gilmour and Mason (joined by Wright) continued recording and touring under the name Pink Floyd. Waters used legal means to try to keep them from using the name, declaring Pink Floyd a spent force, but the parties reached an out-of-court settlement allowing Gilmour, Mason and Wright to continue as Pink Floyd. The band again enjoyed worldwide success with *A Momentary Lapse of Reason* (1987) and *The Division Bell* (1994), and Waters continued as a solo musician, releasing three studio albums. Although for some years relations between Waters and the remaining three members were sour, the band reformed in 2005 for what would be a final one-off performance at Live 8.

Contents [hide]

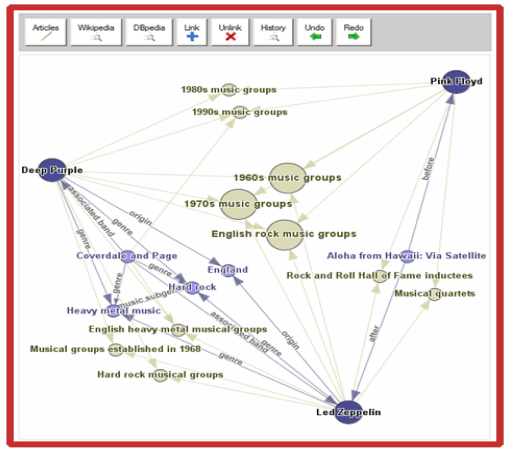

All structured data components (infobox, categories, visualization) are placed in the same vicinity

Pink Floyd

Pink Floyd in 1968 (from left to right): Nick Mason, Syd Barrett, David Gilmour (seated in front), Roger Waters, Richard Wright

Background information

Also known as	The Tea Set, The Pink Floyd Sound, The Pink Floyd
Origin	London
Genres	Progressive rock Psychedelic rock
Years active	1965–1996, 2005
Labels	EMI, Harvest, Capitol, Tower, Columbia
Associated acts	Sigma 6, Joker's Wild, Zee
Website	www.pinkfloyd.co.uk www.pinkfloyd.com

Former members

Syd Barrett
David Gilmour
Bob Klose
Nick Mason
Roger Waters
Richard Wright

Categories: Pink Floyd | Musical groups established in 1965 | Musical groups disestablished in 1996 | English progressive rock groups | Psychedelic rock music groups | 1960s music groups | 1970s music groups | 1980s music groups | 1990s music groups | Musical quartets | Rock and Roll Hall of Fame inductees | Grammy Award winners | Capitol Records artists | Polar Music Prize laureates

History [edit]

Early years (1963–1967) [edit]

Formation [edit]

Nick Mason and Roger Waters met at the Regent Street Polytechnic in London, where both were studying architecture.[1] The pair first played together in a band formed by Keith Noble and Clive Metcalfe, along with Noble's sister Sheilagh, an occasional singer. They were joined later by fellow student Richard Wright. With the addition of Wright the band became a sextet, and took the name Sigma 6.[2] Wright's girlfriend Juliette Gale was often a guest artist, and Waters initially played rhythm guitar before switching to bass. Early gigs were for private functions, and the band rehearsed in a tearoom in the basement of Regent Street Polytechnic. Sigma 6 played songs by The Searchers as well as material written by fellow student Ken Chapman, who became their manager and songwriter.[2]

SmallWorlds

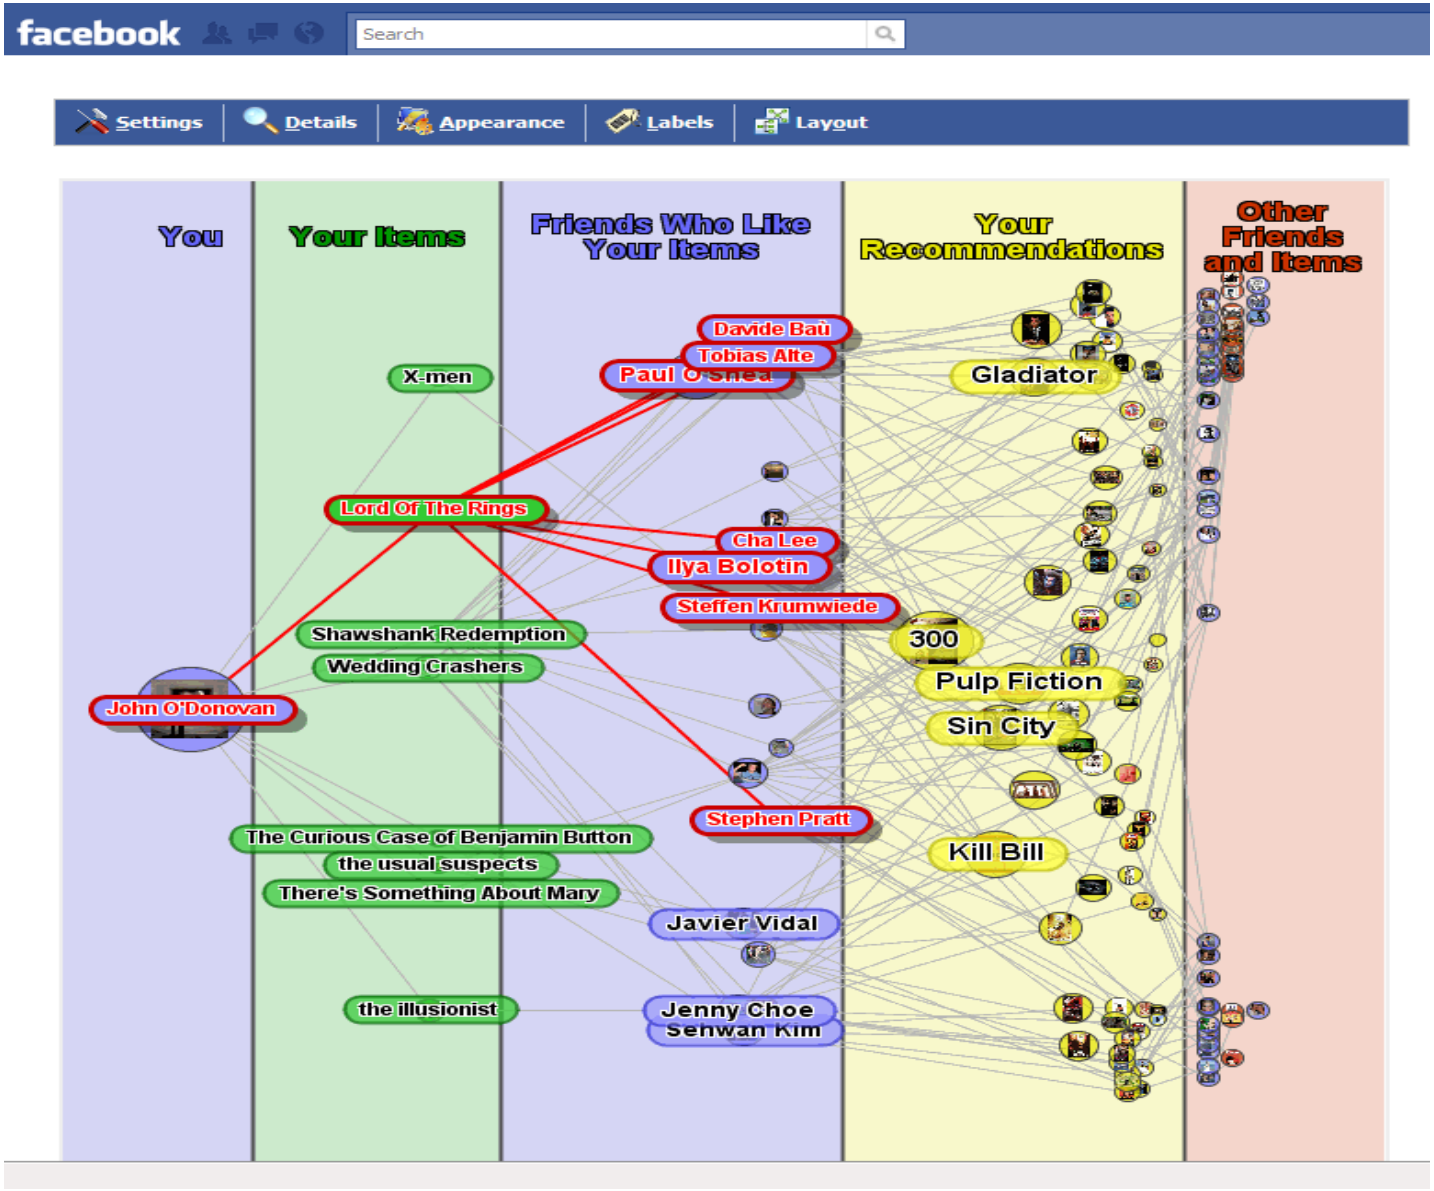

TopicNets

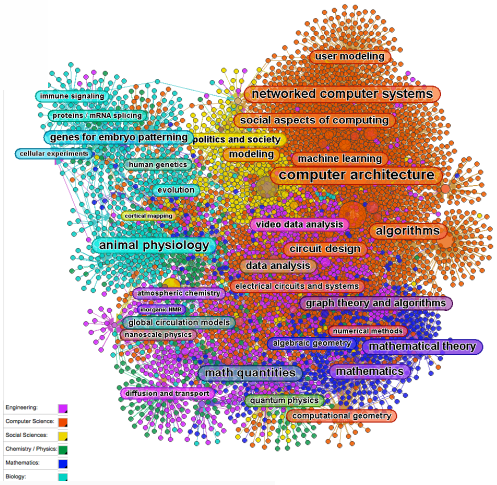

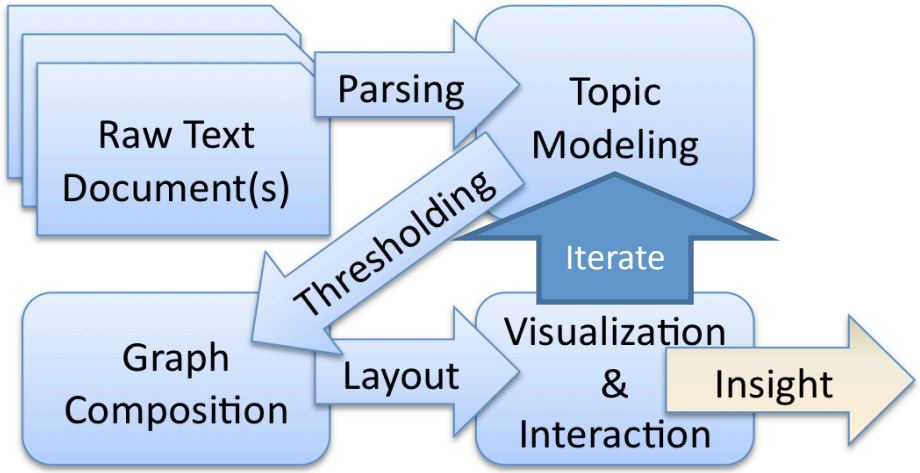

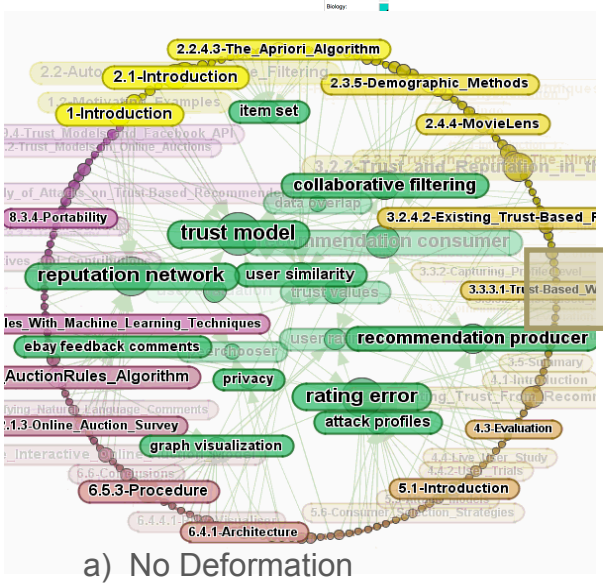

a) No Deformation

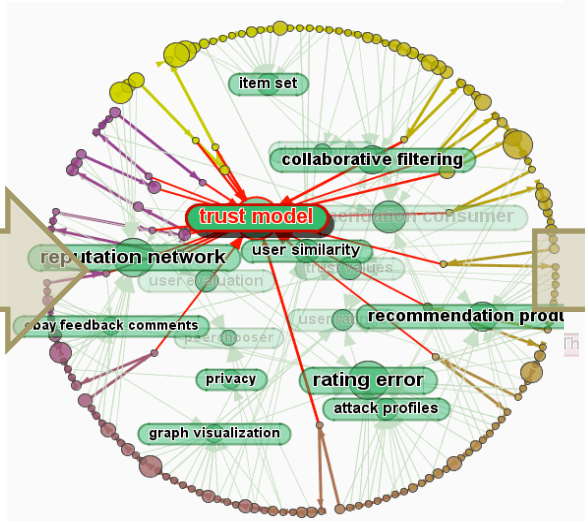

b) Single-Topic Deformation

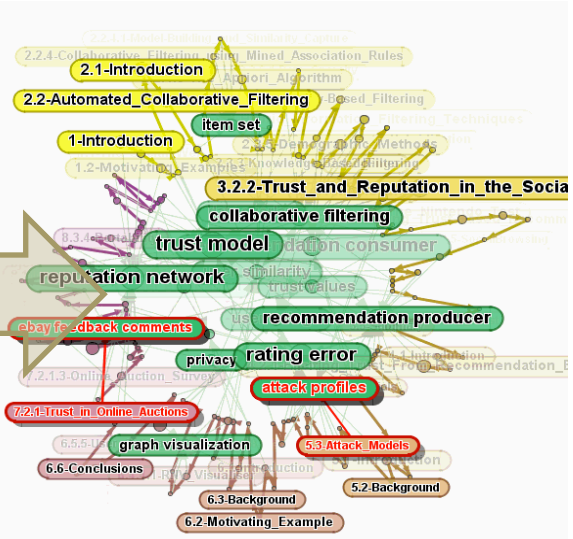

c) All-Topic Deformation

Some new directions for this work:

- Trust and Provenance API for WiGis Framework. … allowing different trust models to plug in seamlessly, so that network effects can be visualized and the model can be understood more easily.
- Promotion of this interaction framework within the community of trust researchers.
- TopicTrust: Using LDA techniques to infer trust between different entities though their topic relations. (using our novel and fast algorithms from TopicNets)

Research Questions

- What are the roles of both interface and interaction in trust-based systems.
- How does transparency introduced by the interface change the user's experience.
- Does transparency and interaction (dynamic provision of data) effect robustness and stability of a trust based system?
- How can Linked Data and better semantic relations be leveraged to build a reliable, open and portable trust model?

http://inismor.ucd.ie/~johno/amazon_mockup/

amazon.com | John's Amazon.com | See All 41 Product Categories | Your Account | Cart | Your Lists | Help | NEW

Search Amazon.com GO | Find Gifts | Web Search GO

Thanks for your order!

We'll send you an e-mail confirmation shortly.

Note: If you ordered several items to be delivered to the same address, we may send them to you in separate boxes to give you the speediest service. Rest assured, this will not affect your shipping charges.

Manage Your Order

- Track the status of this order
- Cancel items from this order
- Edit Shipping Method
- See the status of all your orders

Recommended for You

The Amityville Horror
(Why is this recommended for you?)

Hostel
(Why is this recommended for you?)

The Legend of Zelda: Twilight Princess
(Why is this recommended for you?)

› See more Recommendations

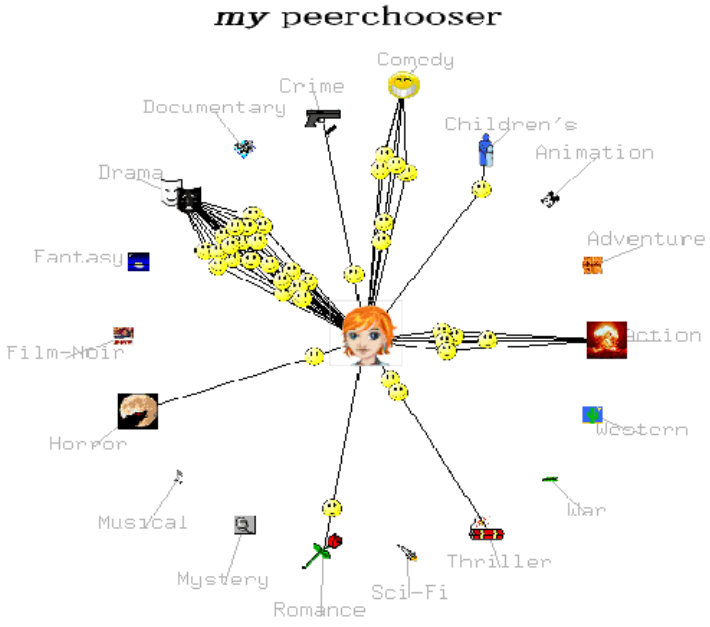

Personalizing Trust in Online Auctions

John O'Donovan [+], Vesile Evrim[*], Barry Smyth[+]. Dennis McLeod[*], Paddy Nixon[+].

(+) Adaptive Information Cluster, Department of Computer Science, University College Dublin
(*) Semantic Information Lab, Viterbi School of Engineering. University of Southern California

(Publications: STAIRS. Riva Del Garda AUG 28th 2006, IJCAI 2007, Hyderabad, India)

Motivations

•Buyers and sellers in online auctions are faced with the task of deciding who to entrust their business to based on a very limited amount of information.

•Over 99% of the feedback comments left on eBay are positive. This is a large bias and can be misleading.

Motivations

- There is a large amount of *hidden* negativity in. feedback comments

- For Example: User (*a)* gives a positive rating for a transaction with user (b) and leaves the following: "Nice Item, but shipping was too expensive and the item was poorly packaged"

- This scenario occurs frequently on auction sites such as eBay because users are afraid to leave negative comments for fear of retaliatory negative feedback. (Resnick 02)

Main Point

•We introduce a new way to calculate trust for online auctions based on classification of user feedback and show how the resultant trust can be tailored for different users.

Background

Gamon et al. [3] introduce the concept of varying the level of granularity during the classification procedure, finding that "varying the level of granularity of analysis allows the discovery of new information".

Yukari et al. (‘04) Introduces the idea of salient *satisfaction factors* as a mechanism for classifying negative comments.

•Pang and Lee (‘02) Experiments with sentiment classification in freetext comments. Found that machine learning algorithms generally performed poorly at this task.

Extracting Trust from Feedback Comments

Sub-Problem: automatically attaining accurate classification of feedback comments into +'ve and –'ve

Solutions: Machine Learning?

(no specific domain knowledge needed, scalable, portable to different auctions, training data…, what type of ML algorithm is best? How accurate/reliable??)

Natural Language Processing?

(domain knowledge required- but how much?, is it portable to other auctions without changing the domain knowledge? How accurate?)

NLP Approach: AuctionRules Algorithm

We designed an NLP classification algorithm specifically for extracting the polarity of sentiment in online auction comments…

Representing the Auction as a Set of Features

1. *Item* - The quality/condition of the product being bought or sold. (*item, product*)
2. *Person* - The person the user makes the transaction with. (*buyer, seller, eBayer*, *dealer*)
3. *Cost* - Cost of item, cost of shipping, hidden costs etc. (*expense, cost*)
4. *Shipping* - Delivery of the item, security, time etc. (*delivery, shipping*)
5. *Response* - Communication with the other party, emails, feedback comment responses. (*response, comment, email, communication*)
6. *Packaging* - The packaging quality/condition of the item *packaging*
7. *Payment* - how the payment will be made to the seller, or back to buyer for return (*payment*)
8. *Transaction* - the overall transaction quality (*Service, Transaction, Business*)

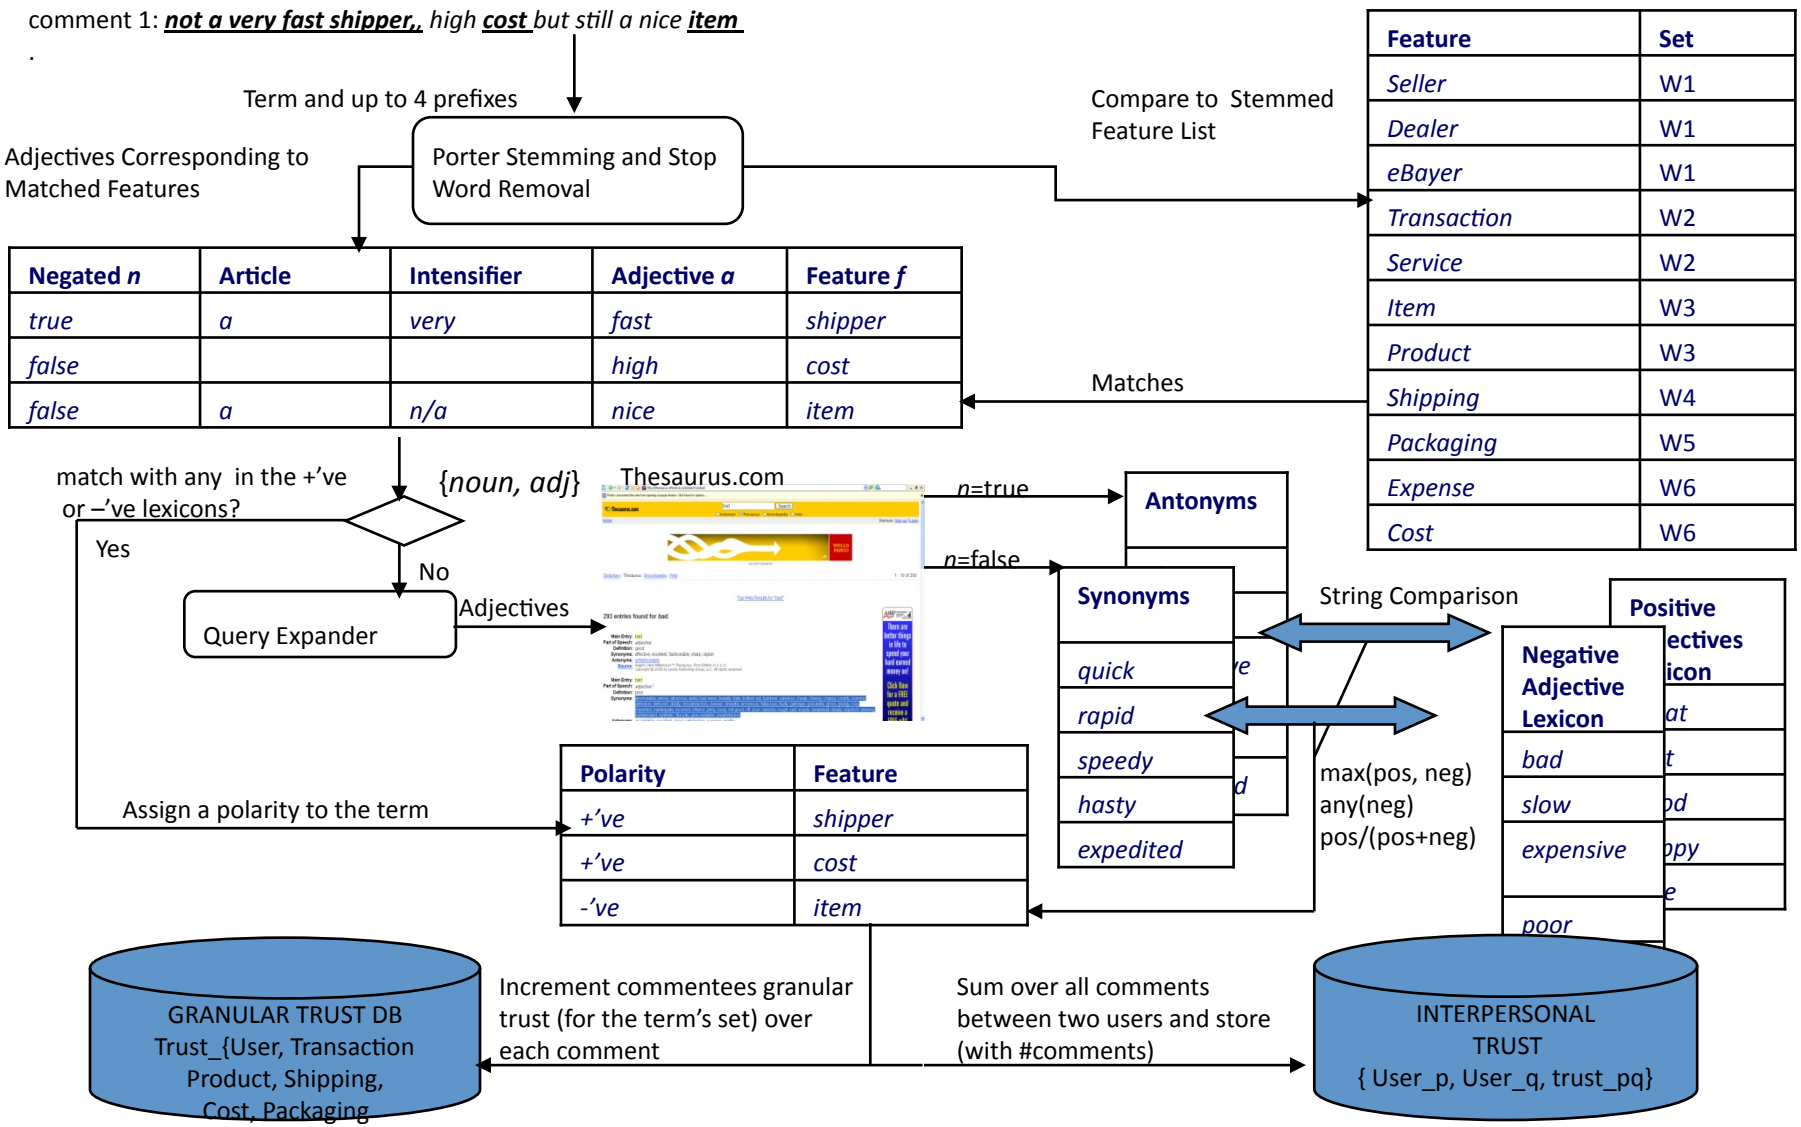
comment 1: not a very fast shipper,, high cost but still a nice item
.
Term and up to 4 prefixes
Porter Stemming and Stop Word Removal
Adjectives Corresponding to Matched Features
Compare to Stemmed Feature List
Feature
Set
Seller
W1
Dealer
W1
eBayer
W1
Transaction
W2
Service
W2
Item
W3
Product
W3
Shipping
W4
Packaging
W5
Expense
W6
Cost
W6
Negated n
Article
Intensifier
Adjective a
Feature f
true
a
very
fast
shipper
false
high
cost
false
a
n/a
nice
item
Matches
match with any in the +'ve or –'ve lexicons?
{noun, adj}
Thesaurus.com
n=true
n=false
Antonyms
Synonyms
quick
rapid
speedy
hasty
expedited
String Comparison
Yes
No
Query Expander
Adjectives
Positive Adjectives Lexicon
Negative Adjective Lexicon
bad
slow
expensive
poor
max(pos, neg)
any(neg)
pos/(pos+neg)
Assign a polarity to the term
Polarity
Feature
+'ve
shipper
+'ve
cost
-'ve
item
GRANULAR TRUST DB
Trust_{User, Transaction Product, Shipping, Cost, Packaging
Increment commentees granular trust (for the term's set) over each comment
Sum over all comments between two users and store (with #comments)
INTERPERSONAL TRUST
{ User_p, User_q, trust_pq}

A Social Network Diagram

eBay Trust Network

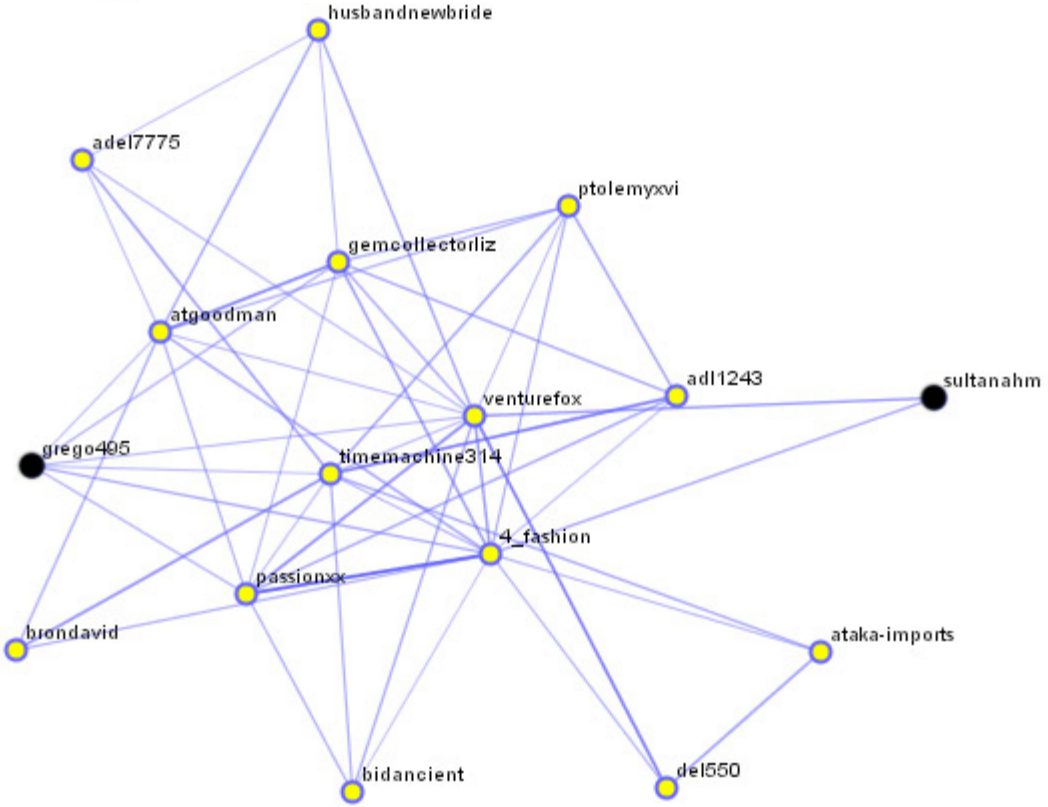

Generated by TrustMe- Graph Generator on johnod.net; using PieSpy 0.4.0
Blue edge thickness and shortness represents strength of relationship
http://www.jibble.org/piespy/ - This frame was drawn at Mon May 08 23:26:23 PDT 2006

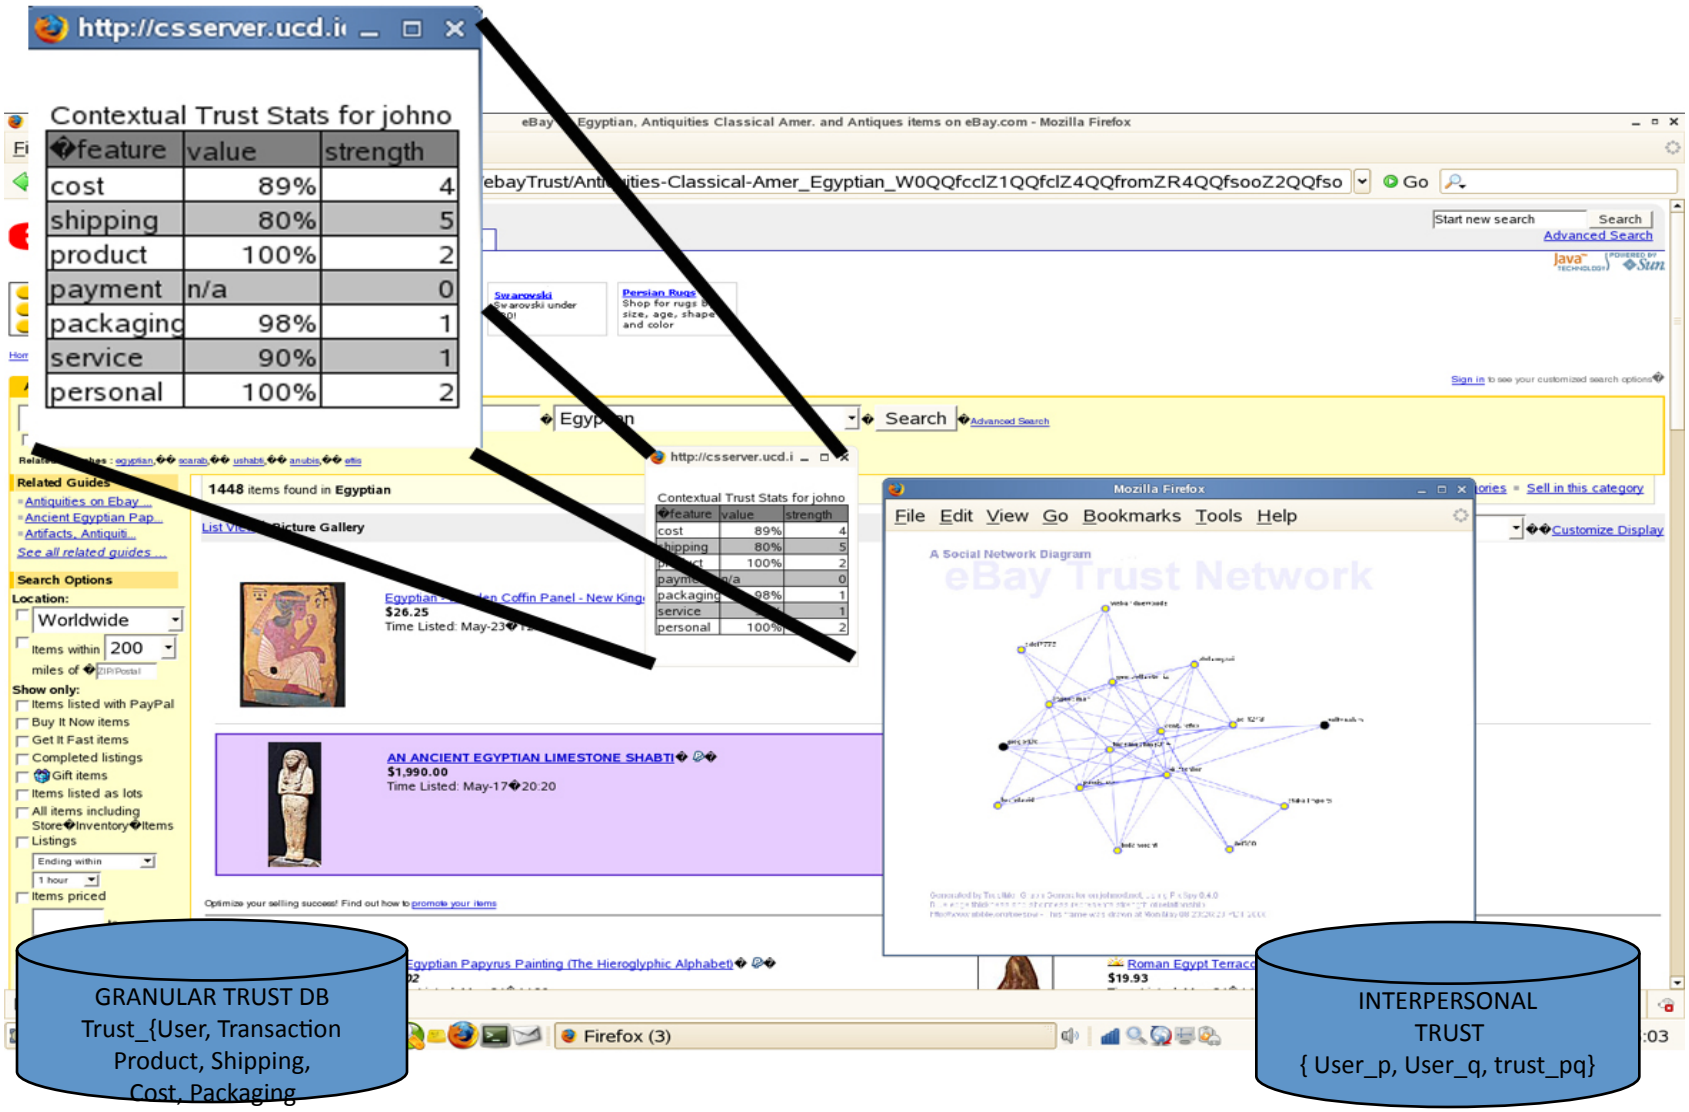
http://csserver.ucd.ie
Contextual Trust Stats for johno
feature | value | strength
cost | 89% | 4
shipping | 80% | 5
product | 100% | 2
payment | n/a | 0
packaging | 98% | 1
service | 90% | 1
personal | 100% | 2
eBay Egyptian, Antiquities Classical Amer. and Antiques items on eBay.com - Mozilla Firefox
Mozilla Firefox
File Edit View Go Bookmarks Tools Help
A Social Network Diagram
eBay Trust Network
1448 items found in Egyptian
AN ANCIENT EGYPTIAN LIMESTONE SHABTI
$1,990.00
Time Listed: May-17 20:20
Egyptian Papyrus Painting (The Hieroglyphic Alphabet)
Roman Egypt Terrac
$19.93
Firefox (3)
GRANULAR TRUST DB
Trust_{User, Transaction
Product, Shipping,
Cost, Packaging
INTERPERSONAL
TRUST
{ User_p, User_q, trust_pq}

Trust and Incentives

A Dynamical Perspective

Munindar P. Singh
(Joint work with Christopher J. Hazard)
singh@ncsu.edu

North Carolina State University

Motivating Question and Findings

What are the consequences of approaching trust from the standpoint of incentives?

- Natural axioms relating trust and incentives
- Trustworthiness maps to patience (discount factor)
- Reputation models as mechanisms
  - Govern agents' behaviors
  - Can be analyzed as dynamical systems
- Intuitive desiderata for reputation models map to technical properties
- Yield a principled basis for comparing existing systems
  - Contrast with traditional, anecdotal evaluations

Defining Trustworthiness

Scalar	
Preference	Prefer b to c if b would fulfill more costly commitment than c
Strength	Do something costly $\Rightarrow$ will do something cheaper
Stability	Preferences stable if time shifted

Intertemporal Discounting and Its Estimation

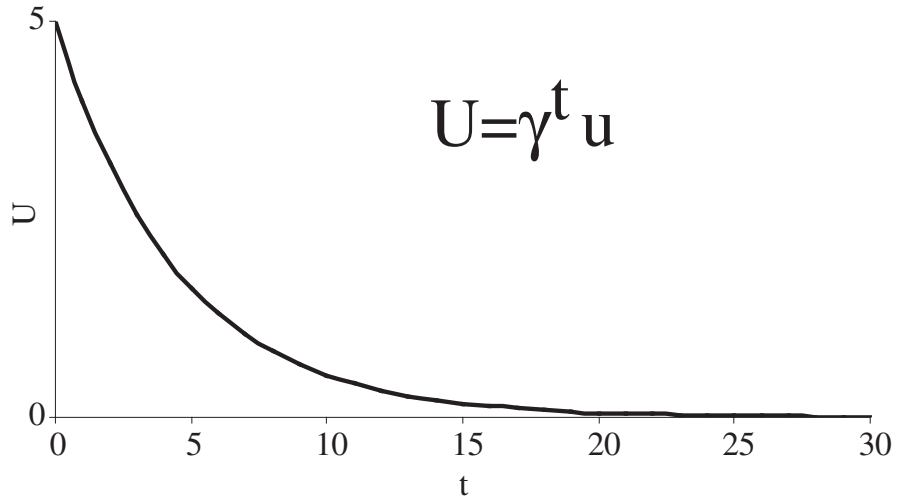

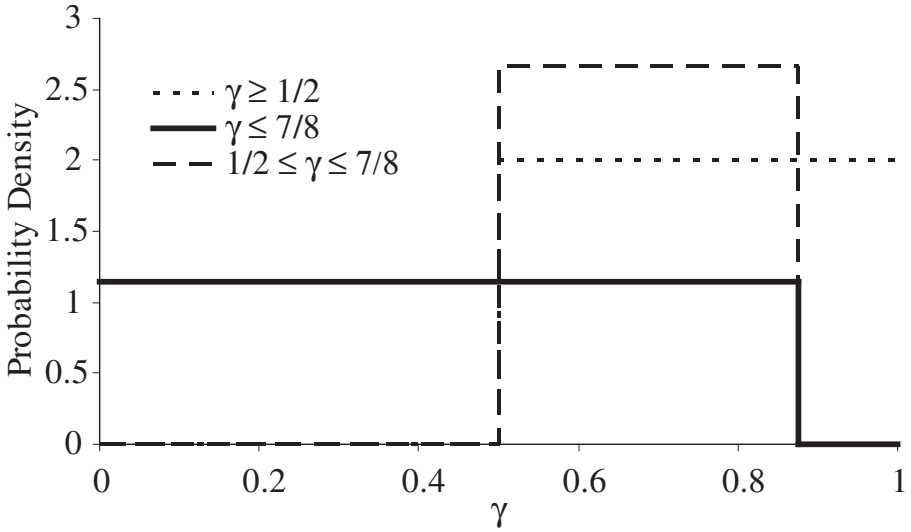

Trustworthiness and Discount Factor Isomorphism

Proof Sketch

- Compare two agents
- From Preference and Scalar properties, examine one event
- Individual rationality $\Rightarrow$ agent expects a net positive utility (future utility from relationship)
- If b will fulfill commitment that c will not, b's perceived utility $>$ c's perceived utility
- Stability: perceived utility inequalities $\Rightarrow$ constant multiplicative factor, x^t (because $x^{t+s} = x^t \cdot x^s$)
- x can break inequalities from trustworthiness definition $\Rightarrow$ x related to trustworthiness
- Utility and domain of x match standard discounting

Trust and Reputation Dynamics

- Agents' utilities are functions of their reputations
- Reputation dynamics with strategic agent
  - Formalizing key desiderata
- Conceptual interface for reputation systems
  - *Update function:* next reputation after action
  - *Payoff function:* reward for action given reputation
- Compare reputation systems

Trust System Metrics: Notation

- Agent type: $\theta \in \Theta$
- Current reputation (projection): $r \in R$
- Next reputation function: Ω
  - $r' = \Omega_\theta(r)$
- Fixed point reputation function: χ
  - $\chi(\theta) = \lim_{n \to \infty} \Omega_\theta^n(r_{\text{initial}})$

Dynamic Reputation Graphs

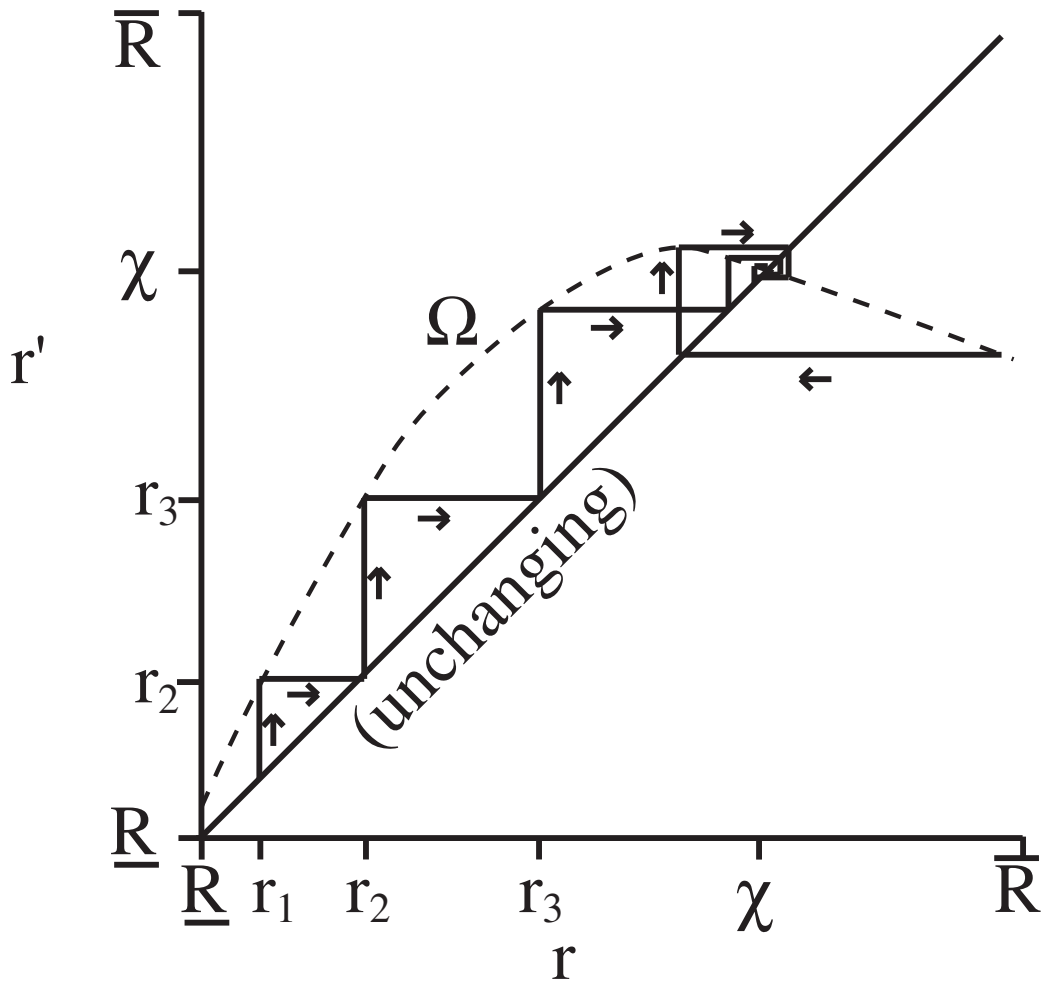

Ideal Trust System

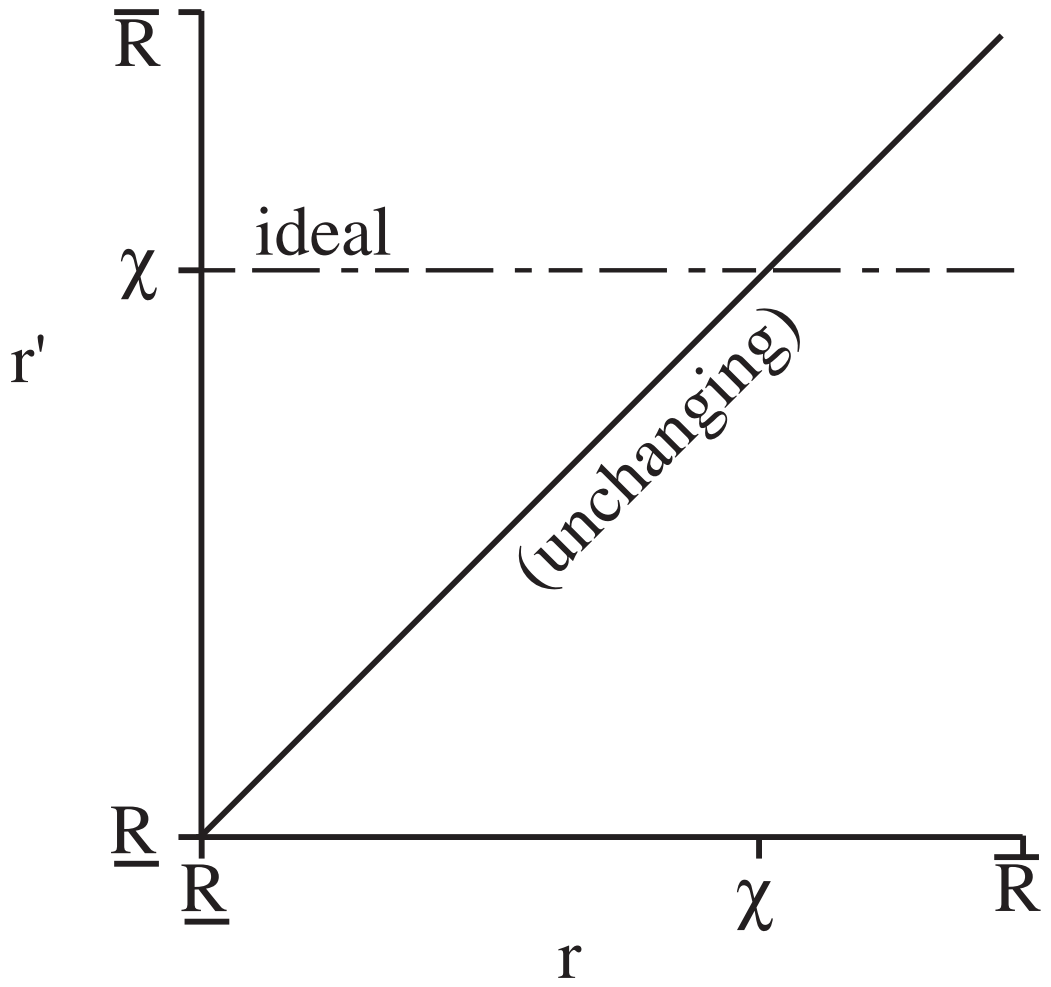

Good Trust System

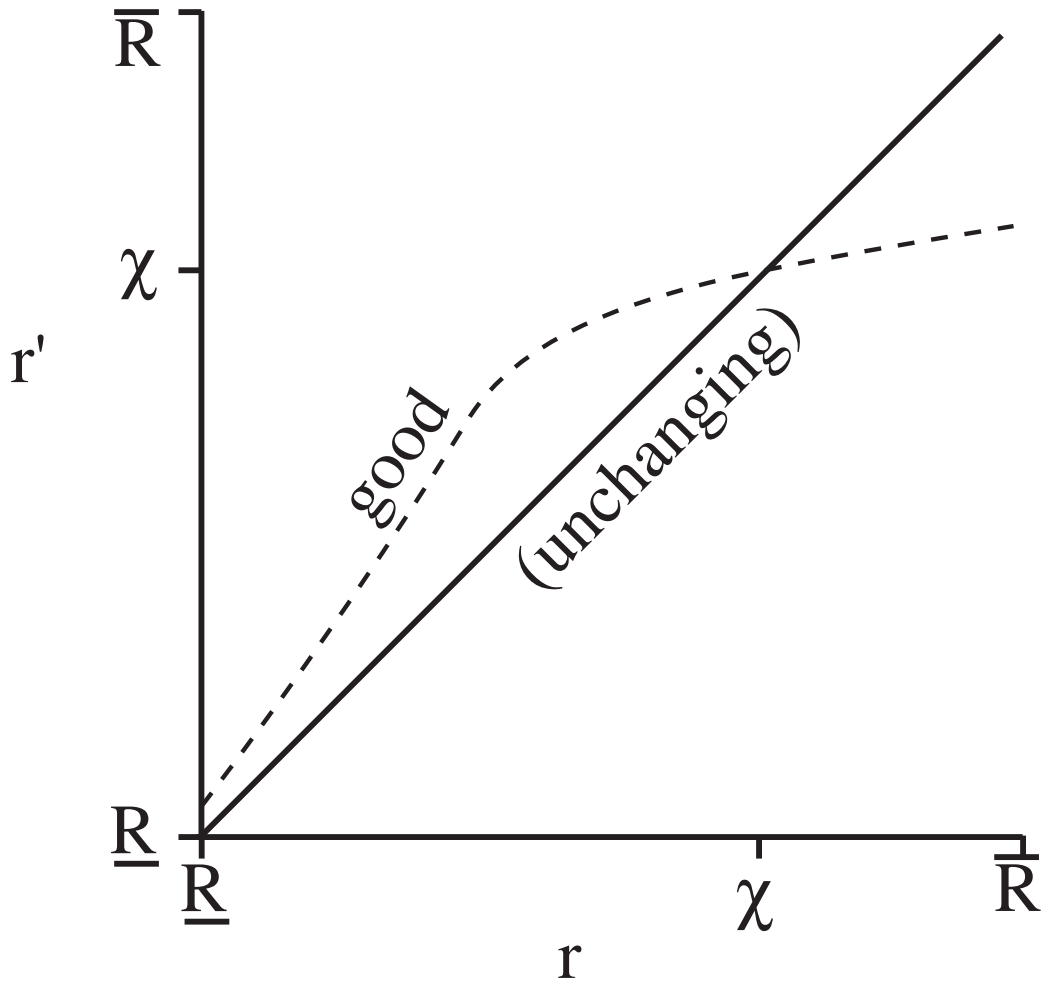

Trust System Metric 1: Unambiguity

Each agent type asymptotically maps to a single reputation value

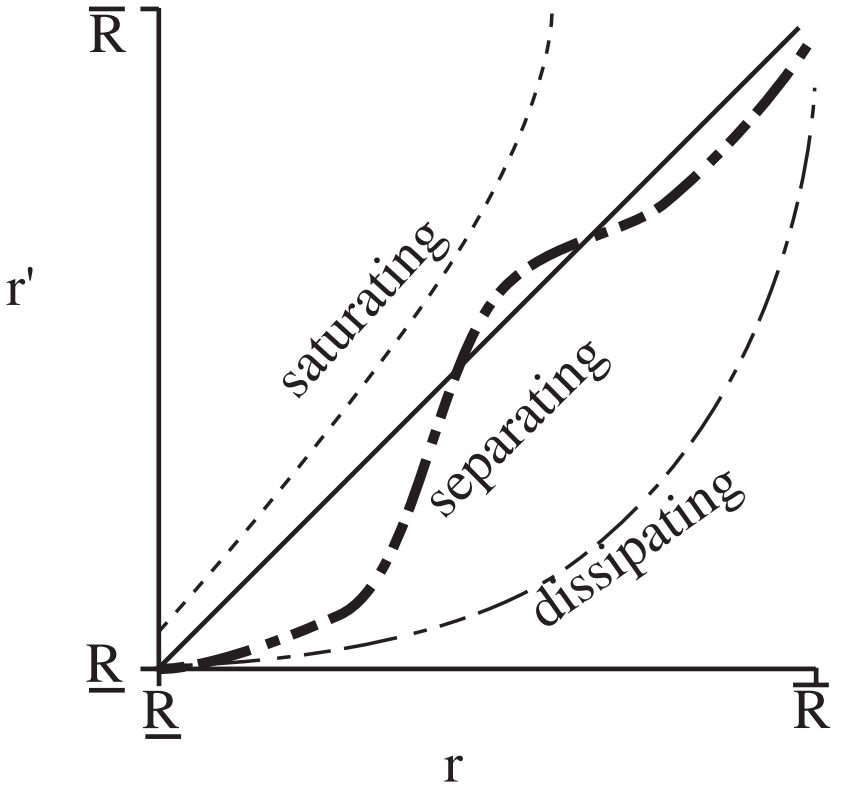

Ambiguous Trust Systems

Trust System Metric 2: Monotonicity

- *Ideally Patient Strategic* (IPS) agent
  - Infinite horizon, maximize utility
- If θ_a is weakly preferable to θ_b to IPS agent $\Rightarrow$ asymptotically, a's reputation $\geq$ b's reputation

Trust System Metric 3: Convergence

Reputation should converge quickly near the fixed point

- Max component of gradient: $||\nabla\Omega(r)||_\infty < 1$ and minimized
- Lyapunov stability may be acceptable

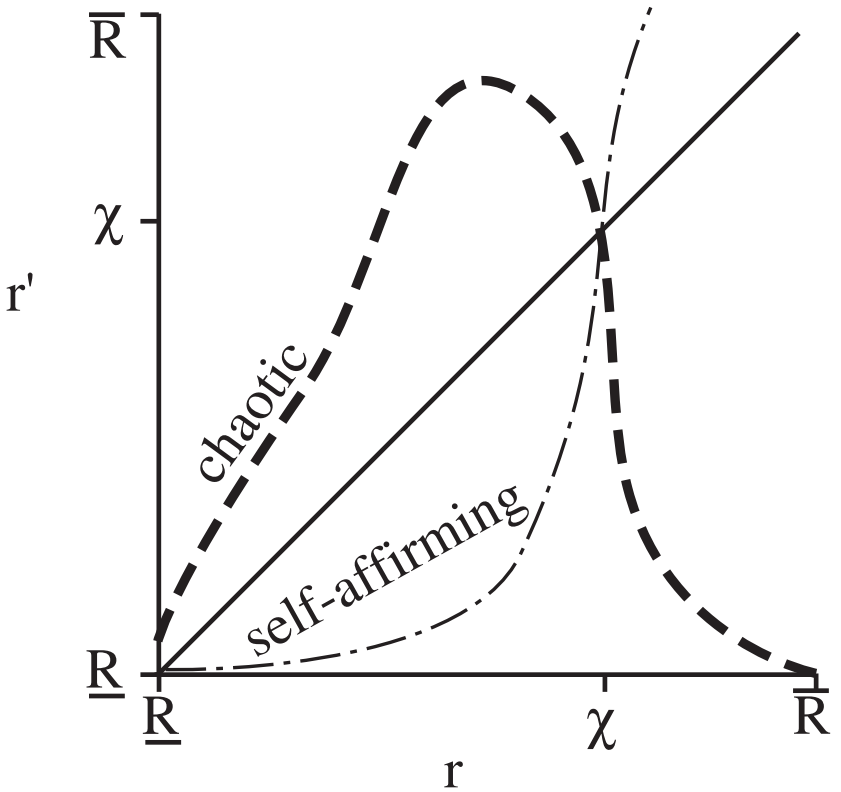

Divergent Trust Systems

Trust System Metric 4: Accuracy

Minimize average error over distribution of types

- Error: absolute distance from ideal reputation
- Reputation system performance when beliefs far from fixed point

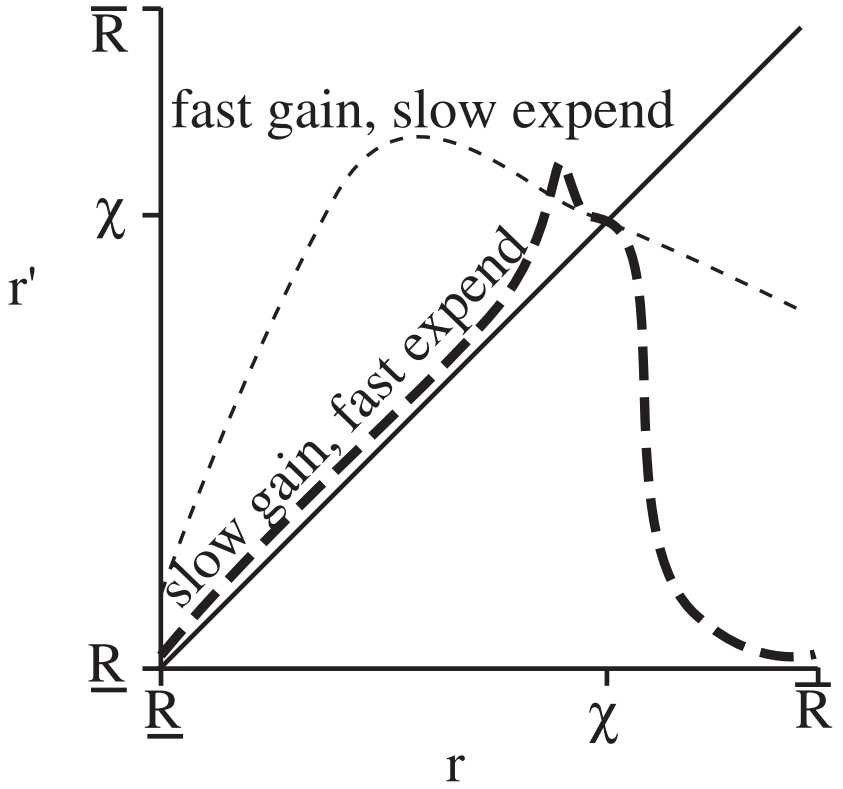

Differing Accuracy

Probabilistic Reciprocity Graph

Probability of favor based on balance [Sen 2002]

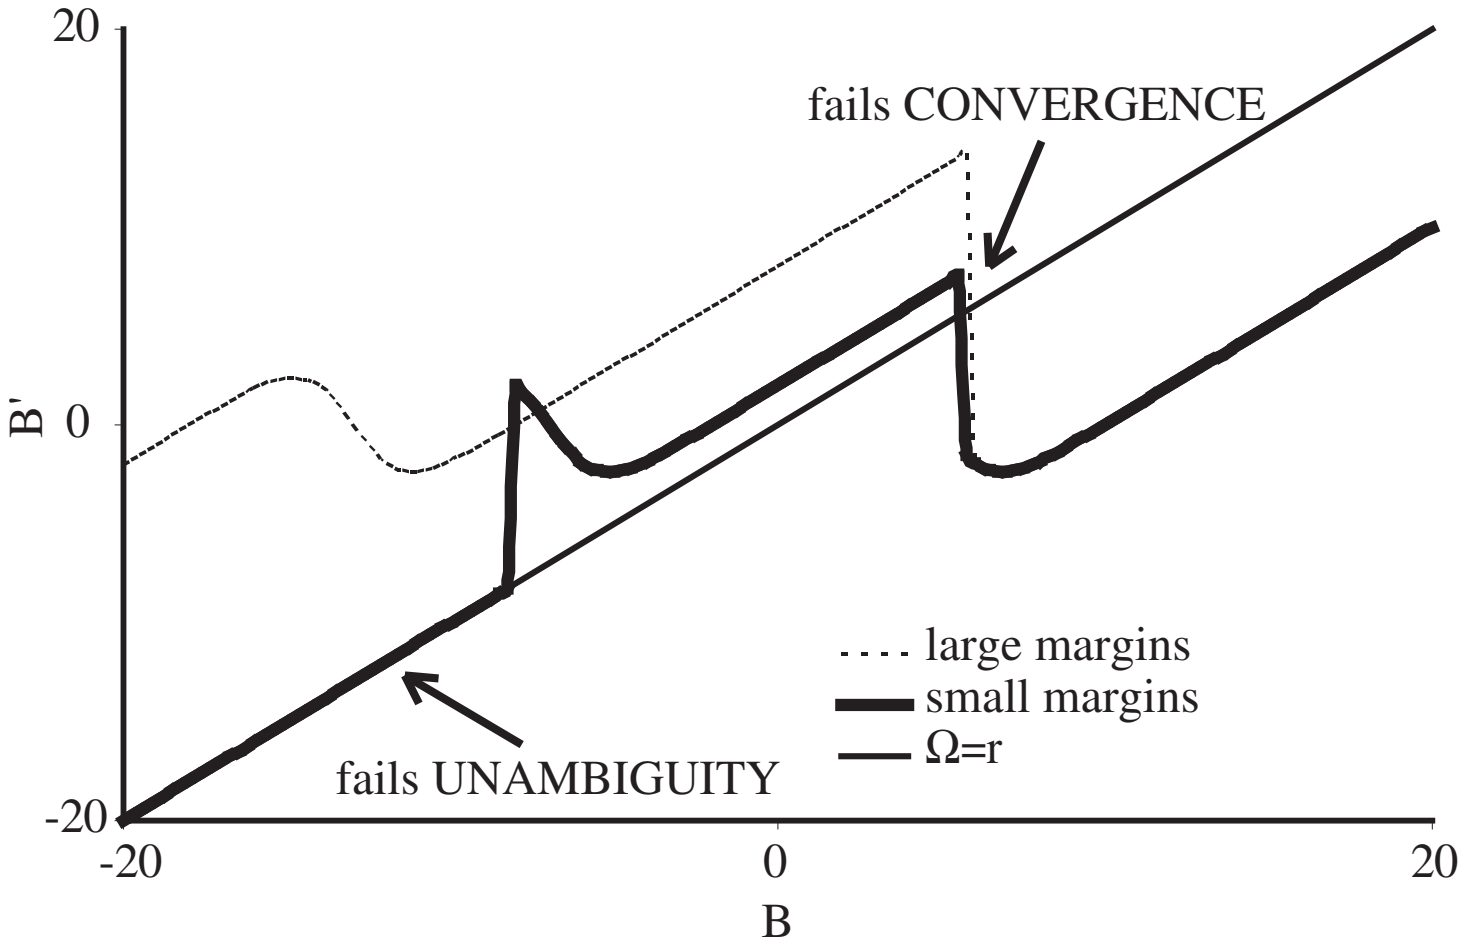

Discount Factor Graph

Trustworthiness $\sim$ patience [Hazard & Singh 2010; Smith & desJardins 2009]

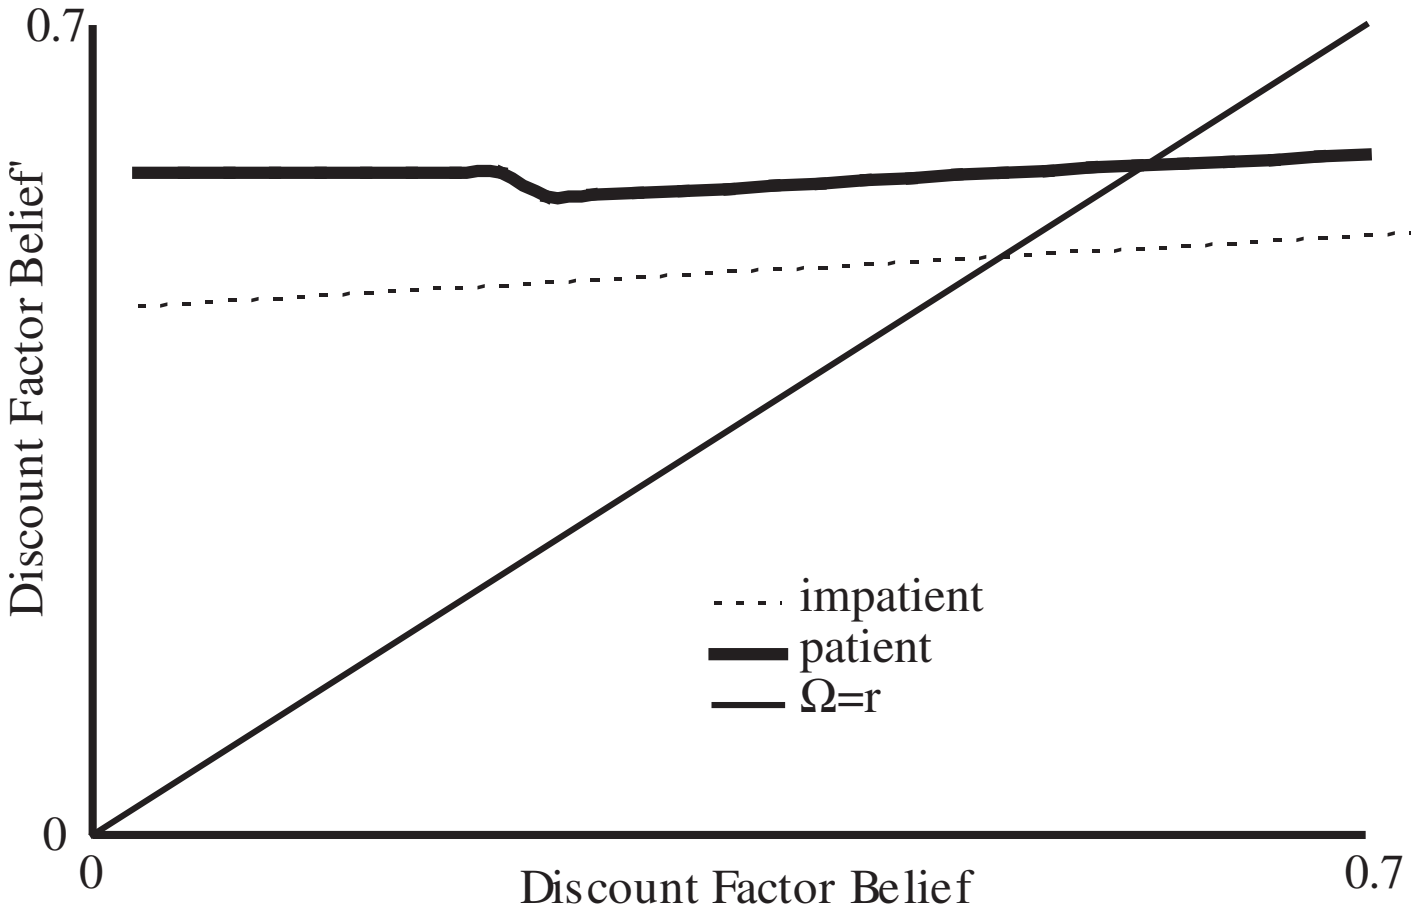

Results

Rep. System	Unambig.	Monotonicity	Convergence	Accuracy
Beta (suplin)	yes	monotonic	0 and 0.9	0.4
Beta (linear, sublin)	yes	nondiscriminatory	0.9	0.45
Certainty	no	–	1	–
Discount Factor	yes	monotonic	< 0.1	0.02
Prob. Reciprocity	no	monotonic	no	0.2
Sporas (suplin, linear)	yes	monotonic	≈ 0	0.3
Sporas (sublin)	yes	nonmonotonic	no	0.4
Travos	yes	monotonic	0.8	0.2

- Beta, Certainty, Travos models [Jøsang 1998; Wang & Singh 2006, 2007; Teacy et al. 2006]
  - Probability of positive and negative interactions
- Sporas [Zacharia & Maes 2000]
  - Reputation measured and dampened

Empirical Reputation Dynamics from Amazon

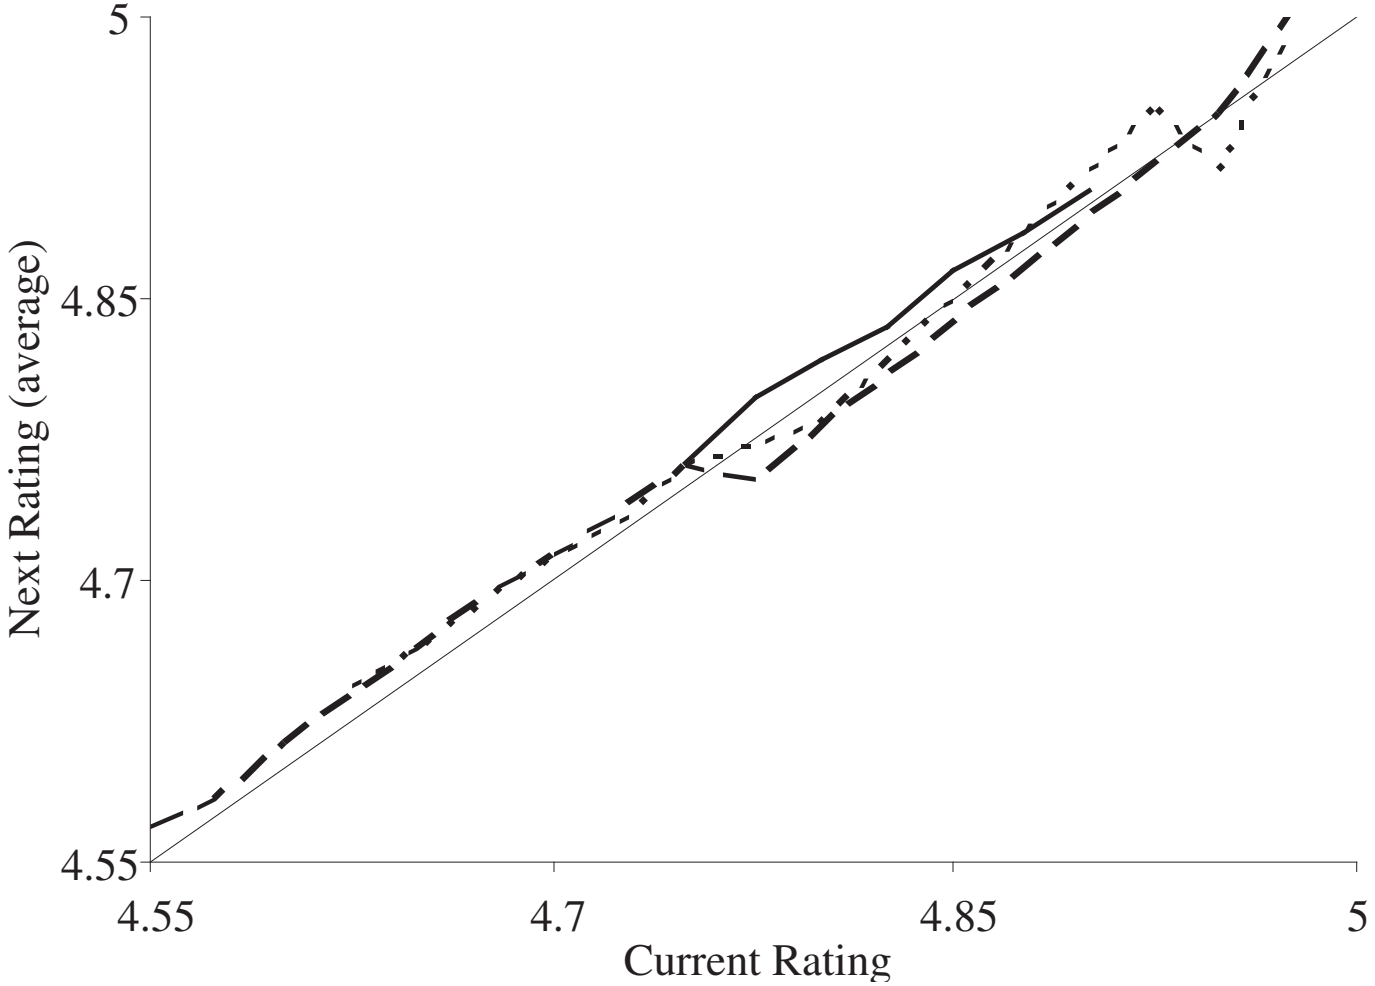

Conclusions and Directions

Beginning from incentives leads naturally to dynamical analyses

- Discount factor isomorphic to trustworthiness given assumptions
- Desiderata are first major step in comparing diverse reputation systems
- Discount factor model performs well, does not yet support continuous choices

Thanks!

http://www.csc.ncsu.edu/faculty/mpsingh/

Trust as a Basis for Social Computing

Munindar P. Singh

North Carolina State University

July 2010

Notions of Trust

Existing literature

- *Subjective*
  - As a conglomerate of mental attitudes
- *Social*
  - Based on social relationships
- *Distributed*
  - Based on certificate chains
- *Measured* applied to the above
  - Based on heuristics, probabilities, utility, . . .

Traditional approaches emphasize estimation over meaning

Applying Trust for Social Computing

Trust underlies all interactions among autonomous parties

- Trust reflects the trustor's *dependence* on the trustee
- Currently, trust is applied
  - Embedded into each specific application
  - Not reusable
- Many types of social relationships, each nuanced
  - Acquaintanceship
  - Friendship
  - Organizational
  - Task-specific
- How may we abstract out trust to apply it as a basis for social computing applications?

Programming Social Applications

Approach

Specify and configure

- Roles
- Social interactions
- Their effects on social states
- Any additional constraints

Programming Social Applications

Middleware

- Offer primitives for social interactions
  - Communicating
  - Maintaining social state
  - Computing trust on behalf of a participant

Social Middleware to Support Social Applications

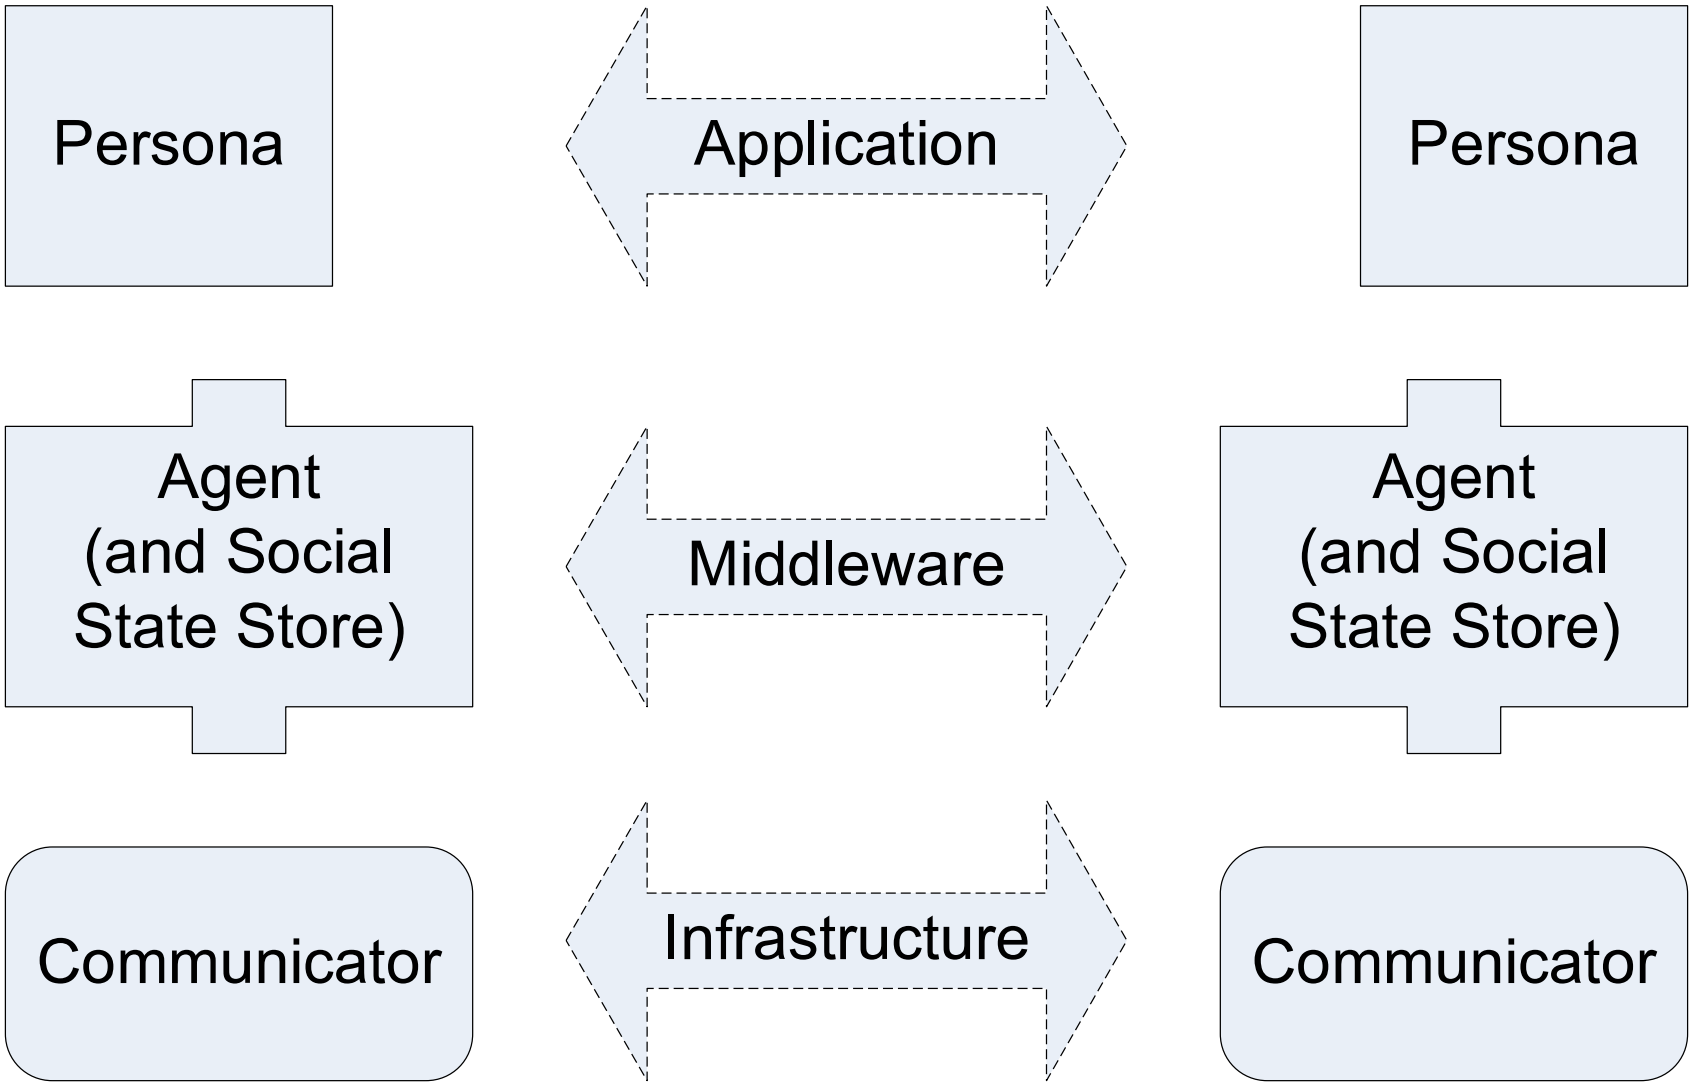

Architecture Conceptually

How a system is organized

- Primarily its ingredients
  - Components
  - Connectors
- An architectural style
  - Constraints on components and connectors
  - Patterns on components and connectors

Architecture Examples

Components; connectors; constraints; patterns

- Electrical system
  - Power elements (sources, sinks); conductors; no short circuits; star
  - How do we characterize the elements and conductors logically?
    - Current drawn, voltage expected, impedance offered
- Social system
  - Individuals; social relationships; ?; ?
  - How do we characterize the individuals and their relationships?
  - *Claim:* Trust is what flows over a relationship
    - Can we characterize relationships in a reusable manner, even though not domain-independent?

Social Middleware Related to Architecture

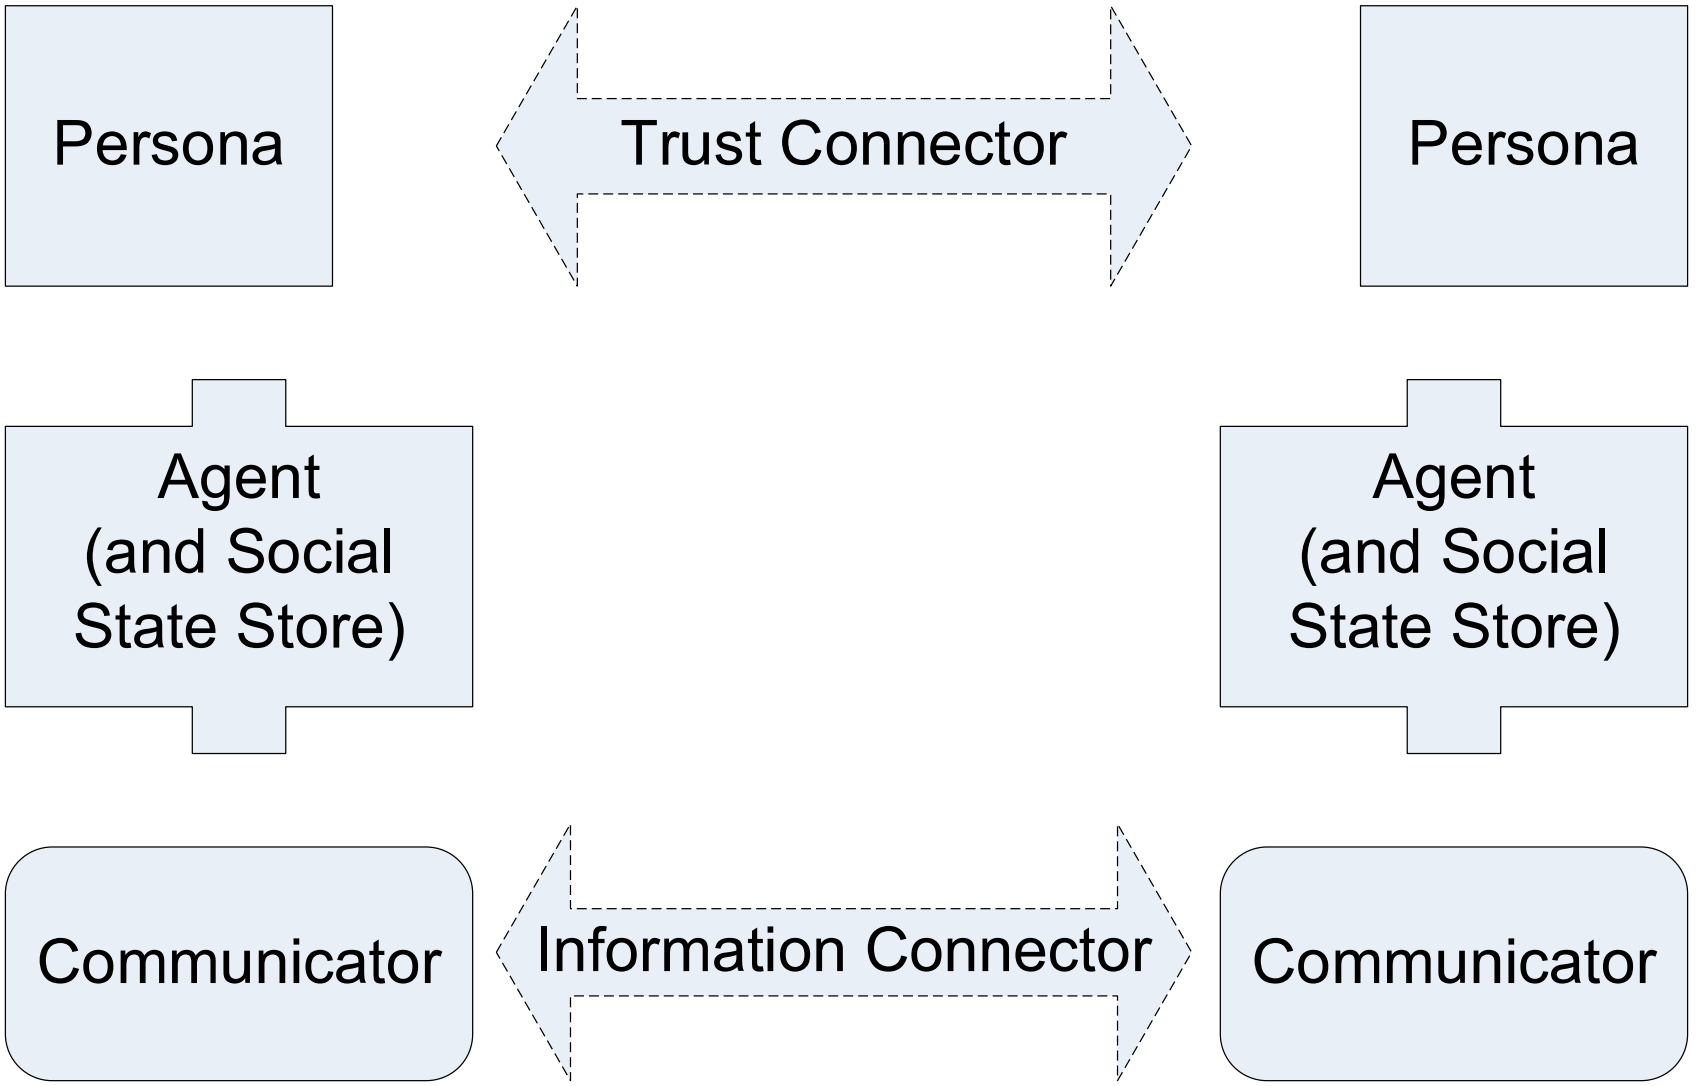

Understanding Trust in Architectural Terms

General Model of Trust

- Notions of dependence
- Conditional
- Compositional
- Semantic
- General

Trust from a Logical Standpoint

- $\mathsf{T}_{trustor,trustee}(antecedent, consequent)$
  - $\mathsf{T}_{Alice,Bob}$(raise alert, send warning)
  - $\mathsf{T}_{trustor,trustee}(\top, consequent)$: unconditional trust
- ACTIVATE: $\mathsf{T}_{x,y}(r,u) \wedge r \rightarrow \mathsf{T}_{x,y}(\top,u)$
  - $\mathsf{T}_{Alice,Bob}$(raise alert, send warning) $\wedge$ raise alert $\Rightarrow \mathsf{T}_{Alice,Bob}(\top$, send warning)
- COMPLETE: $u \rightarrow \neg\mathsf{T}_{x,y}(r,u)$
  - send warning $\Rightarrow \neg\mathsf{T}_{Alice,Bob}$(raise alert, send warning)
  - send warning $\Rightarrow \neg\mathsf{T}_{Alice,Bob}(\top$, send warning)

A formal semantics underlies the above notion

Schematic of an Architectural Connector as Trust

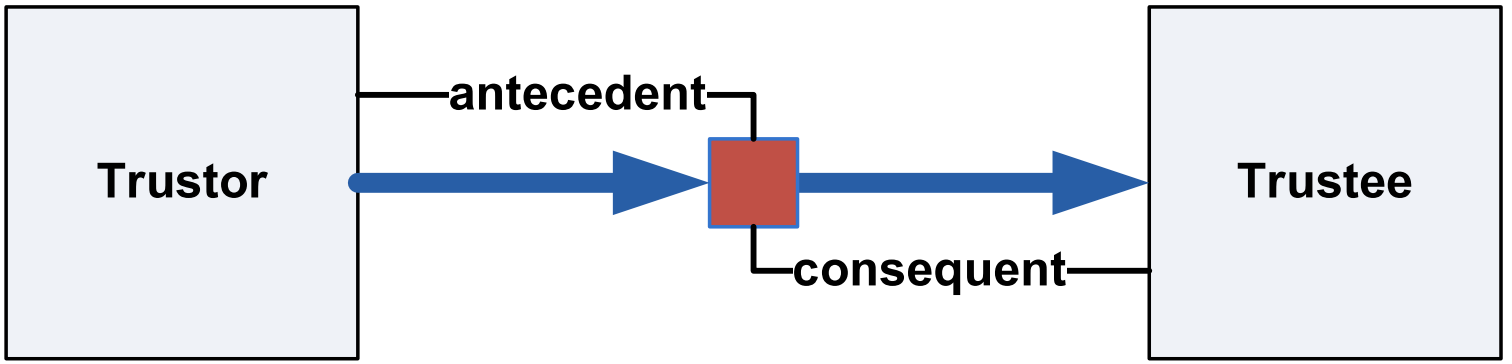

Postulates for Trust

Active trust basics

(Omitting *trustor* and *trustee* when they are the same throughout)

- Complete a connector: dependence has been fulfilled
  - $u \rightarrow \neg \mathsf{T}(r, u)$
- Activate a connector: make dependence stronger (strongest when $r = \top$)
  - $\mathsf{T}(r \wedge s, u) \wedge r \rightarrow \mathsf{T}(s, u)$
- Partition a connector: a dependence for two things is a dependence for each separately (if it isn't already done)
  - $\mathsf{T}(r, u \wedge v) \wedge \neg u \rightarrow \mathsf{T}(r, u)$

Postulates for Trust

Connector integrity

- Avoid conflict: dependence must be internally consistent
  - $\mathsf{T}(r, u) \rightarrow \neg \mathsf{T}(r, \neg u)$
- Nonvacuity: dependence must be for something tangible
  - From $r \vdash u$ infer $\neg \mathsf{T}(r, u)$
- Tighten: if a dependence holds then a narrower dependence also holds
  - From $\mathsf{T}(r, u), s \vdash r, s \nvdash u$ infer $\mathsf{T}(s, u)$

Postulates for Trust

Connector structure

- Combine antecedents: two connectors with the same consequent (fulfillment condition) yield a broader connector
  - $\mathsf{T}(r, u) \wedge \mathsf{T}(s, u) \rightarrow \mathsf{T}(r \vee s, u)$
- Combine consequents: two connectors with the same antecedent (trigger condition) yield a stronger connector
  - $\mathsf{T}(r, u) \wedge \mathsf{T}(r, v) \rightarrow \mathsf{T}(r, u \wedge v)$
- Chain: two chained dependencies yield a combined dependence
  - From $\mathsf{T}(r, u), u \vdash s, \mathsf{T}(s, v)$ infer $\mathsf{T}(r, v)$

Postulates for Trust

Connector meaning

- Exposure: the trustee's commitment is its level of exposure if the trustor trusts it for it
  - $C_{x,y}(r,u) \rightarrow T_{y,x}(r,u)$
- Transient alignment: when the trustee commits to support the dependency
  - $T_{x,y}(r,u) \rightarrow C_{y,x}(r,u)$
- Well-placed trust: when trust is fulfilled in the actual execution
  - $T_{x,y}(\text{true},u) \rightarrow Ru$
- Whole-hearted alignment: when trust is backed by a steady commitment until success
  - $T_{x,y}(s,v) \rightarrow R(s \rightarrow (C_{y,x}(s,v) U v))$

(Above, $C_{x,y}(r,u)$ refers to a commitment from x to y; R indicates "on the real execution path"; and pUq means p holds until q does)

TRUSTEE'S TEAM, Schematically

If you trust a team member, you trust the team

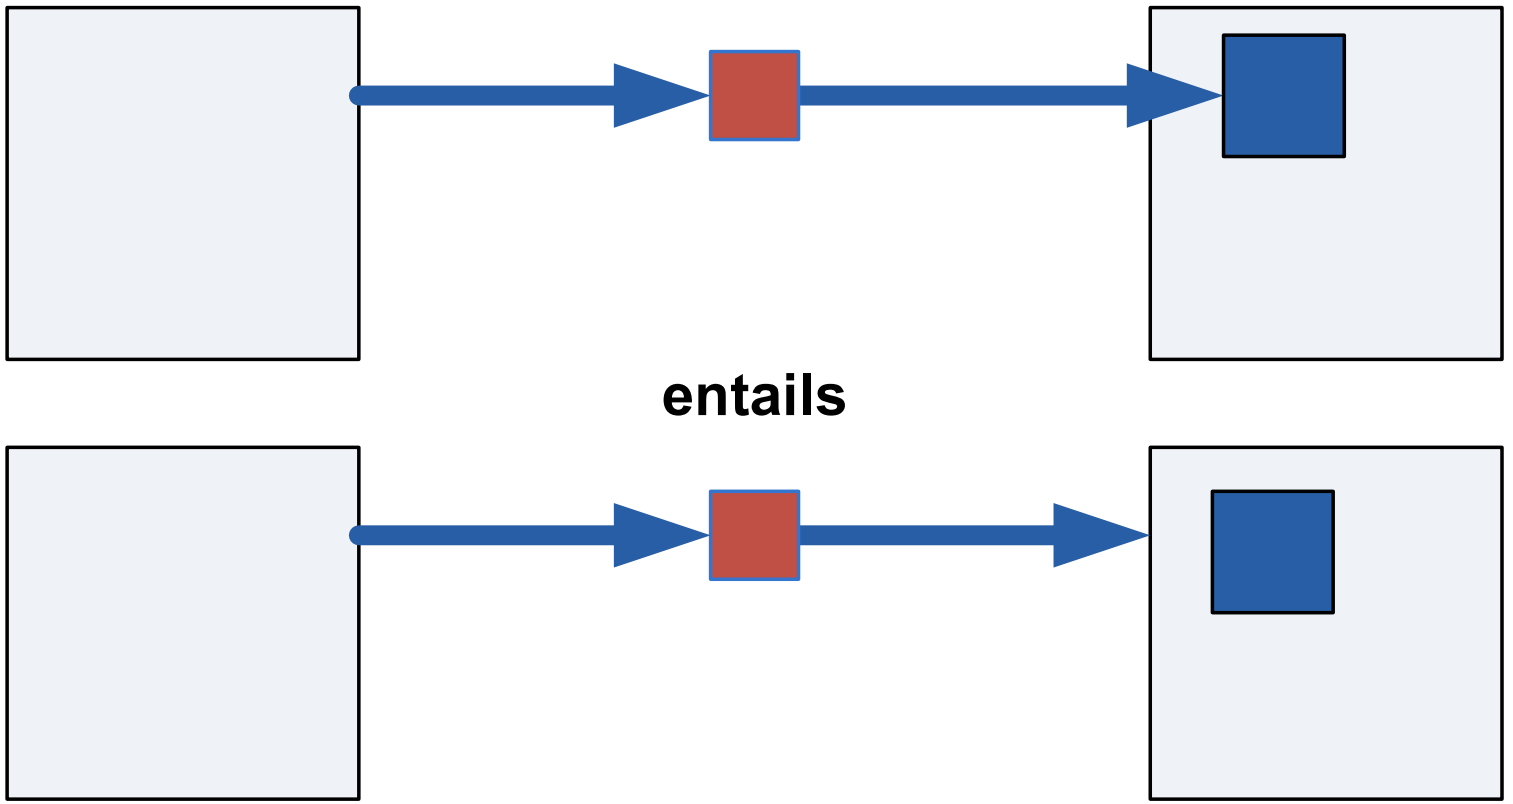

TRANSIENT ALIGNMENT, Schematically

The trustee is committed to what you trust them for

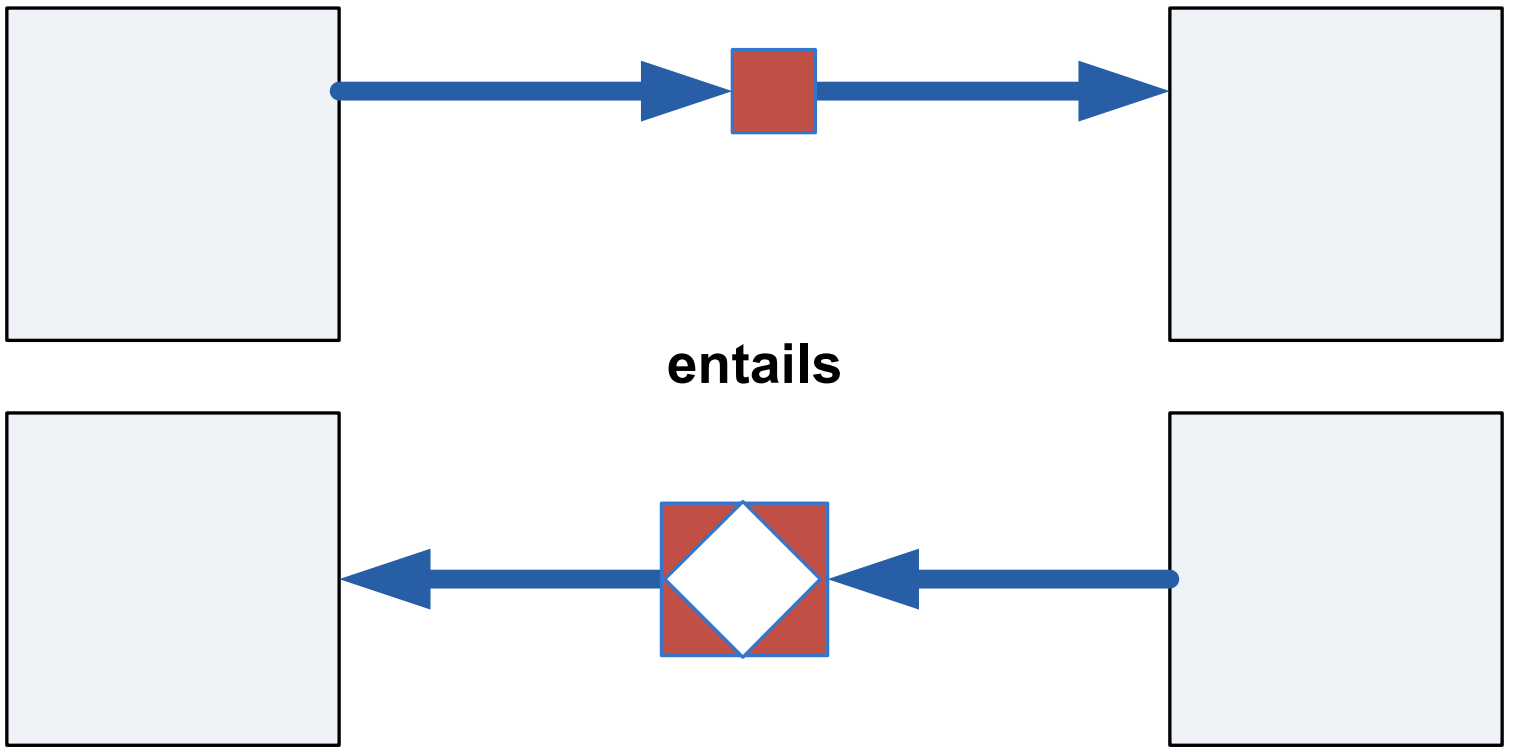

TRUSTOR'S TEAM, Schematically

Your team trusts whom you trust

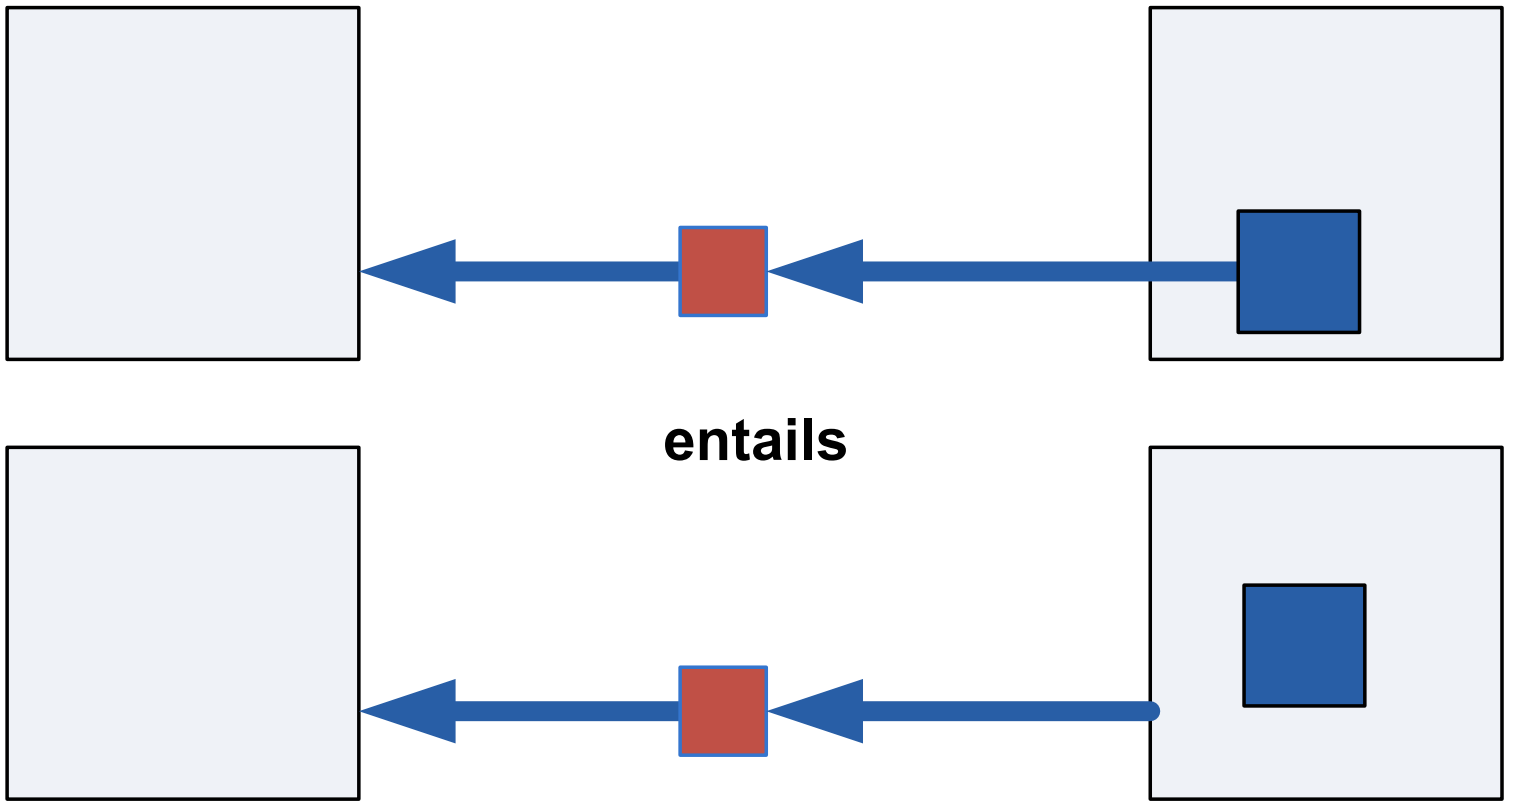

PARALLEL TEAMWORK, Schematically

If you trust each other, you are part of a team

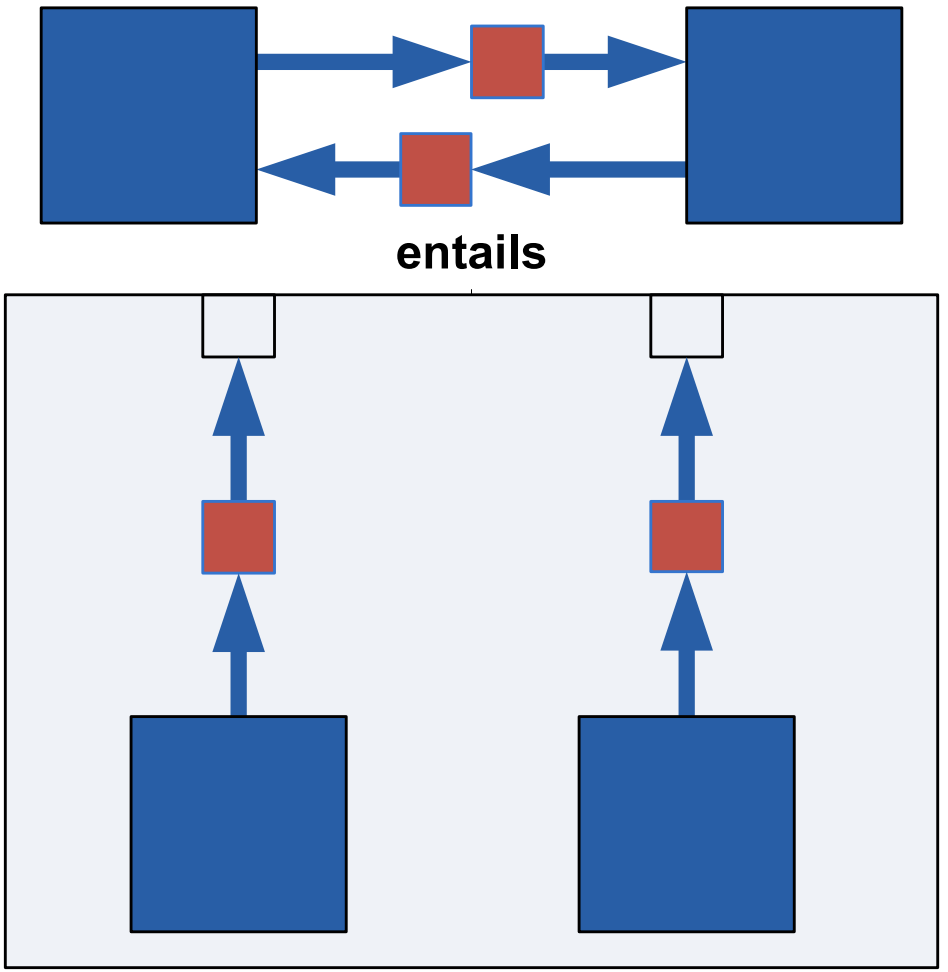

Directions

- Formalizing architectures for social computing based on trust
  - How can trust fit into an overall system architecture?
- Identifying suitable architecture styles
  - What are suitable patterns for different types of social applications?
- Mapping effectively to existing representations and estimation techniques
  - Computation paths can be used as a basis for judging probabilities and expected utilities

Research topics

1) Linguistics translation to machine understanding (Ron, Chris) -2
2) Dynamics of trust (Cho) (Munindar) 4
   - How trust changes over time?
   - Forgiveness
   - Regretfulness
3) Trust/risk model, impact on decision making/uncertainty management (Kevin) 2,5
4) Human behavior to Trust (context, intention, motivation, cooperative vs. self-centered, reward/incentives) (Ugur) 4
5) Trust aggregation and inference (Ron, Ugur) 4
6) Distrust (Jen)1.2
7) Property of trust systems (virtual trust plane) (Felix) 3.5
8) Framework for integrating and leveraging of trust relationship (Mike) 3
9) Engineering of trust systems (incentive) (Munindar) 1
10) Communication of trust (John) -3